

Introduction

1.1 Scope

The scope of IEEE Std. 1003.1-200x is described in the Base Definitions volume of IEEE Std. 1003.1-200x.

1.2 Conformance

Conformance requirements for IEEE Std. 1003.1-200x are defined in the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 2, Conformance.

1.3 Normative References

Normative references for IEEE Std. 1003.1-200x are defined in the Base Definitions volume of IEEE Std. 1003.1-200x.

1.4 Changes from Issue 4

Notes to Reviewers

This section with side shading will not appear in the final copy. - Ed.

The change history is subject to revision. The intent is to document changes from Issue 4 thru Issue 6, with the latest change history also documenting changes from the ISO POSIX-1: 1996 standard.

The following sections describe changes made to this volume of IEEE Std. 1003.1-200x since Issue 4. The CHANGE HISTORY section for each entry details the technical changes that have been made to that entry since Issue 4. Changes made between Issue 2 and Issue 4 are not included.

1.4.1 Changes from Issue 4 to Issue 4, Version 2

The following list summarizes the major changes that were made in this volume of IEEE Std. 1003.1-200x from Issue 4 to Issue 4, Version 2:

- The X/Open UNIX extension was added. This specifies the common core APIs of 4.3 Berkeley Software Distribution (BSD 4.3), the OSF AES, and SVID Issue 3.
- STREAMS were added as part of the X/Open UNIX extension.
- Existing Issue 4 functions were clarified as a result of industry feedback.

28 1.4.2 Changes from Issue 4, Version 2 to Issue 5

29 The following list summarizes the major changes that were made in this volume of
30 IEEE Std. 1003.1-200x from Issue 4, Version 2 to Issue 5:

- 31 • Functions previously defined in the ISO POSIX-2 standard C-language Binding, Shared
32 Memory, Enhanced Internationalization, and X/Open UNIX Extension Feature Groups were
33 moved to the BASE.
- 34 • Threads were added to the BASE for alignment with the POSIX Threads Extension.
- 35 • The Realtime Threads Feature Group was added.
- 36 • The Realtime Feature Group was added for alignment with the POSIX Realtime Extension.
- 37 • Multibyte Support Extensions (MSE) were added to the BASE for alignment with
38 ISO/IEC 9899:1990/Amendment 1:1995 (E).
- 39 • Large File Summit (LFS) Extensions were added to the BASE for support of 64-bit or larger
40 files and file systems.
- 41 • X/Open-specific threads extensions were added to the BASE.
- 42 • X/Open-specific dynamic linking functions were added to the BASE.
- 43 • A new category Legacy was added.
- 44 • The categories TO BE WITHDRAWN and WITHDRAWN were removed.

45 1.4.3 Changes from Issue 5 to Issue 6 (IEEE Std. 1003.1-200x)

46 The following list summarizes the major changes that were made in this volume of
47 IEEE Std. 1003.1-200x from Issue 5 to Issue 6:

- 48 • This volume of IEEE Std. 1003.1-200x is extensively revised so it can be both an IEEE POSIX
49 Standard and an Open Group Technical Standard.
- 50 • The POSIX System Interfaces requirements incorporate support of FIPS 151-2.
- 51 • The POSIX System Interfaces requirements are updated to align with some features of the
52 Single UNIX Specification.
- 53 • A RATIONALE section is added to each reference page.

54 **1.5 New Features**55 **1.5.1 New Features in Issue 4, Version 2**

56 The functions, headers, and external variables first introduced in Issue 4, Version 2 are listed in
 57 the table below.

58

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

New Functions, Headers, and External Variables in Issue 4, Version 2				
FD_CLR()	endutxent()	gettimeofday()	ptsname()	sigaltstack()
FD_ISSET()	expm1()	getutxent()	putmsg()	sighold()
FD_SET()	fattach()	getutxid()	putpmsg()	sigignore()
FD_ZERO()	fchdir()	getutxline()	pututxline()	siginterrupt()
_longjmp()	fchmod()	getwd()	random()	sigpause()
_setjmp()	fchown()	grantpt()	re_comp()	sigrelse()
a64l()	fcvt()	ilogb()	re_exec()	sigset()
acosh()	fdetach()	index()	readlink()	sigstack()
asinh()	ffs()	initsate()	readv()	srandom()
atanh()	fntmsg()	insque()	realpath()	statvfs()
basename()	fstatvfs()	ioctl()	regcmp()	strcasecmp()
bcmp()	ftime()	isastream()	regex()	strdup()
bcopy()	ftok()	killpg()	remainder()	strncasecmp()
brk()	ftruncate()	l64a()	remque()	swapcontext()
bsd_signal()	gcvt()	lchown()	rindex()	symlink()
bzero()	getcontext()	lockf()	rint()	sync()
cbrt()	getdate()	log1p()	sbrk()	syslog()
closelog()	getdtablesize()	logb()	scalb()	tcgetsid()
dbm_clearerr()	getgrent()	lstat()	select()	truncate()
dbm_close()	gethostid()	makecontext()	setcontext()	ttyslot()
dbm_delete()	getitimer()	mknod()	setgrent()	ualarm()
dbm_error()	getmsg()	mkstemp()	setitimer()	unlockpt()
dbm_fetch()	getpagesize()	mktemp()	setlogmask()	usleep()
dbm_firstkey()	getpgid()	mmap()	setpgrp()	utimes()
dbm_nextkey()	getpmsg()	mprotect()	setpriority()	valloc()
dbm_open()	getpriority()	msync()	setpwent()	vfork()
dbm_store()	getpwent()	munmap()	setregid()	wait3()
dirname()	getrlimit()	nextafter()	setreuid()	waitid()
ecvt()	getrusage()	nftw()	setrlimit()	writev()
endgrent()	getsid()	openlog()	setstate()	
endpwent()	getsubopt()	poll()	setutxent()	
<fmtmsg.h>	<re_comp.h>	<sys/resource.h>	<sys/uio.h>	<utmpx.h>
<libgen.h>	<strings.h>	<sys/statvfs.h>	<sys/un.h>	
<ndbm.h>	<stropts.h>	<sys/time.h>	<syslog.h>	
<poll.h>	<sys/mman.h>	<sys/timeb.h>	<ucontext.h>	
getdate_err	__loc1			

96 **1.5.2 New Features in Issue 5**

97 The functions and headers first introduced in Issue 5 are listed in the table below.

	New Functions and Headers in Issue 5		
100	<i>aio_cancel()</i>	<i>pthread_attr_getstacksize()</i>	<i>pthread_self()</i>
101	<i>aio_error()</i>	<i>pthread_attr_init()</i>	<i>pthread_setcancelstate()</i>
102	<i>aio_fsync()</i>	<i>pthread_attr_setdetachstate()</i>	<i>pthread_setcanceltype()</i>
103	<i>aio_read()</i>	<i>pthread_attr_setguardsize()</i>	<i>pthread_setconcurrency()</i>
104	<i>aio_return()</i>	<i>pthread_attr_setinheritsched()</i>	<i>pthread_setschedparam()</i>
105	<i>aio_suspend()</i>	<i>pthread_attr_setschedparam()</i>	<i>pthread_setspecific()</i>
106	<i>aio_write()</i>	<i>pthread_attr_setschedpolicy()</i>	<i>pthread_sigmask()</i>
107	<i>asctime_r()</i>	<i>pthread_attr_setscope()</i>	<i>pthread_testcancel()</i>
108	<i>btowc()</i>	<i>pthread_attr_setstackaddr()</i>	<i>putc_unlocked()</i>
109	<i>clock_getres()</i>	<i>pthread_attr_setstacksize()</i>	<i>putchar_unlocked()</i>
110	<i>clock_gettime()</i>	<i>pthread_cancel()</i>	<i>pwrite()</i>
111	<i>clock_settime()</i>	<i>pthread_cleanup_pop()</i>	<i>rand_r()</i>
112	<i>ctime_r()</i>	<i>pthread_cleanup_push()</i>	<i>readdir_r()</i>
113	<i>dlclose()</i>	<i>pthread_cond_broadcast()</i>	<i>sched_get_priority_max()</i>
114	<i>dLError()</i>	<i>pthread_cond_destroy()</i>	<i>sched_get_priority_min()</i>
115	<i>dlopen()</i>	<i>pthread_cond_init()</i>	<i>sched_getparam()</i>
116	<i>dlsym()</i>	<i>pthread_cond_signal()</i>	<i>sched_getscheduler()</i>
117	<i>fdatasync()</i>	<i>pthread_cond_timedwait()</i>	<i>sched_rr_get_interval()</i>
118	<i>flockfile()</i>	<i>pthread_cond_wait()</i>	<i>sched_setparam()</i>
119	<i>fseeko()</i>	<i>pthread_condattr_destroy()</i>	<i>sched_setscheduler()</i>
120	<i>ftello()</i>	<i>pthread_condattr_getpshared()</i>	<i>sched_yield()</i>
121	<i>ftrylockfile()</i>	<i>pthread_condattr_init()</i>	<i>sem_close()</i>
122	<i>funlockfile()</i>	<i>pthread_condattr_setpshared()</i>	<i>sem_destroy()</i>
123	<i>fwide()</i>	<i>pthread_create()</i>	<i>sem_getvalue()</i>
124	<i>fwprintf()</i>	<i>pthread_detach()</i>	<i>sem_init()</i>
125	<i>fwscanf()</i>	<i>pthread_equal()</i>	<i>sem_open()</i>
126	<i>getc_unlocked()</i>	<i>pthread_exit()</i>	<i>sem_post()</i>
127	<i>getchar_unlocked()</i>	<i>pthread_getconcurrency()</i>	<i>sem_trywait()</i>
128	<i>getgrgid_r()</i>	<i>pthread_getschedparam()</i>	<i>sem_unlink()</i>
129	<i>getgrnam_r()</i>	<i>pthread_getspecific()</i>	<i>sem_wait()</i>
130	<i>getlogin_r()</i>	<i>pthread_join()</i>	<i>shm_open()</i>
131	<i>getpwnam_r()</i>	<i>pthread_key_create()</i>	<i>shm_unlink()</i>
132	<i>getpwuid_r()</i>	<i>pthread_key_delete()</i>	<i>sigqueue()</i>
133	<i>gmtime_r()</i>	<i>pthread_kill()</i>	<i>sigtimedwait()</i>
134	<i>lio_listio()</i>	<i>pthread_mutex_destroy()</i>	<i>sigwait()</i>
135	<i>localtime_r()</i>	<i>pthread_mutex_getprioceiling()</i>	<i>sigwaitinfo()</i>
136	<i>mbrlen()</i>	<i>pthread_mutex_init()</i>	<i>snprintf()</i>
137	<i>mbrtowc()</i>	<i>pthread_mutex_lock()</i>	<i>strtok_r()</i>
138	<i>mbsinit()</i>	<i>pthread_mutex_setprioceiling()</i>	<i>swprintf()</i>
139	<i>mbsrtowcs()</i>	<i>pthread_mutex_trylock()</i>	<i>swscanf()</i>
140	<i>mlock()</i>	<i>pthread_mutex_unlock()</i>	<i>timer_create()</i>
141	<i>mlockall()</i>	<i>pthread_mutexattr_destroy()</i>	<i>timer_delete()</i>
142	<i>mq_close()</i>	<i>pthread_mutexattr_getprioceiling()</i>	<i>timer_getoverrun()</i>
143	<i>mq_getattr()</i>	<i>pthread_mutexattr_getprotocol()</i>	<i>timer_gettime()</i>
144	<i>mq_notify()</i>	<i>pthread_mutexattr_getpshared()</i>	<i>timer_settime()</i>
145	<i>mq_open()</i>	<i>pthread_mutexattr_gettype()</i>	<i>towctrans()</i>

146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167

New Functions and Headers in Issue 5

<i>mq_receive()</i>	<i>pthread_mutexattr_init()</i>	<i>ttyname_r()</i>
<i>mq_send()</i>	<i>pthread_mutexattr_setprioceiling()</i>	<i>vfwprintf()</i>
<i>mq_setattr()</i>	<i>pthread_mutexattr_setprotocol()</i>	<i>vsnprintf()</i>
<i>mq_unlink()</i>	<i>pthread_mutexattr_setpshared()</i>	<i>vswprintf()</i>
<i>munlock()</i>	<i>pthread_mutexattr_settype()</i>	<i>vwprintf()</i>
<i>munlockall()</i>	<i>pthread_once()</i>	<i>wcrtomb()</i>
<i>nanosleep()</i>	<i>pthread_rwlock_destroy()</i>	<i>wcsrtombs()</i>
<i>pread()</i>	<i>pthread_rwlock_init()</i>	<i>wcsstr()</i>
<i>pthread_atfork()</i>	<i>pthread_rwlock_rdlock()</i>	<i>wctob()</i>
<i>pthread_attr_destroy()</i>	<i>pthread_rwlock_tryrdlock()</i>	<i>wctrans()</i>
<i>pthread_attr_getdetachstate()</i>	<i>pthread_rwlock_trywrlock()</i>	<i>wmemchr()</i>
<i>pthread_attr_getguardsize()</i>	<i>pthread_rwlock_unlock()</i>	<i>wmemcmp()</i>
<i>pthread_attr_getinheritsched()</i>	<i>pthread_rwlock_wrlock()</i>	<i>wmemcpy()</i>
<i>pthread_attr_getschedparam()</i>	<i>pthread_rwlockattr_destroy()</i>	<i>wmemmove()</i>
<i>pthread_attr_getschedpolicy()</i>	<i>pthread_rwlockattr_getpshared()</i>	<i>wmemset()</i>
<i>pthread_attr_getscope()</i>	<i>pthread_rwlockattr_init()</i>	<i>wprintf()</i>
<i>pthread_attr_getstackaddr()</i>	<i>pthread_rwlockattr_setpshared()</i>	<i>wscanf()</i>
<aio.h>	<iso646.h>	<sched.h>
<dlfcn.h>	<mqueue.h>	<semaphore.h>
<inttypes.h>	<pthread.h>	<wctype.h>

168 **1.5.3 New Features in Issue 6**169 ***Notes to Reviewers***170 *This section with side shading will not appear in the final copy. - Ed.*171 A table listing new functions, headers, etc. since the ISO POSIX-1:1996 standard will be added
172 here in a future draft.

173 1.6 Terminology

174 This section appears in the Base Definitions volume of IEEE Std. 1003.1-200x, but is repeated
175 here for convenience:

176 For the purposes of IEEE Std. 1003.1-200x, the following terminology definitions apply:

177 **can**

178 Describes a permissible optional feature or behavior available to the user or application. The
179 feature or behavior is mandatory for an implementation that conforms to
180 IEEE Std. 1003.1-200x. An application can rely on the existence of the feature or behavior.

181 **implementation-defined**

182 Describes a value or behavior that is not defined by IEEE Std. 1003.1-200x but is selected by
183 an implementor. The value or behavior may vary among implementations that conform to
184 IEEE Std. 1003.1-200x. An application should not rely on the existence of the value or
185 behavior. An application that relies on such a value or behavior cannot be assured to be
186 portable across conforming implementations.

187 The implementor shall document such a value or behavior so that it can be used correctly
188 by an application.

189 **legacy**

190 Describes a feature or behavior that is being retained for compatibility with older
191 applications, but which has limitations which make it inappropriate for developing portable
192 applications. New applications should use alternative means of obtaining equivalent
193 functionality.

194 **may**

195 Describes a feature or behavior that is optional for an implementation that conforms to
196 IEEE Std. 1003.1-200x. An application should not rely on the existence of the feature or
197 behavior. An application that relies on such a feature or behavior cannot be assured to be
198 portable across conforming implementations.

199 To avoid ambiguity, the opposite of *may* is expressed as *need not*, instead of *may not*.

200 **shall**

201 For an implementation that conforms to IEEE Std. 1003.1-200x, describes a feature or
202 behavior that is mandatory. An application can rely on the existence of the feature or
203 behavior.

204 For an application or user, describes a behavior that is mandatory.

205 **should**

206 For an implementation that conforms to IEEE Std. 1003.1-200x, describes a feature or
207 behavior that is recommended but not mandatory. An application should not rely on the
208 existence of the feature or behavior. An application that relies on such a feature or behavior
209 cannot be assured to be portable across conforming implementations.

210 For an application, describes a feature or behavior that is recommended programming
211 practice for optimum portability.

212 **undefined**

213 Describes the nature of a value or behavior not defined by IEEE Std. 1003.1-200x which
214 results from use of an invalid program construct or invalid data input.

215 The value or behavior may vary among implementations that conform to
216 IEEE Std. 1003.1-200x. An application should not rely on the existence or validity of the
217 value or behavior. An application that relies on any particular value or behavior cannot be

218 assured to be portable across conforming implementations.

219 **unspecified**

220 Describes the nature of a value or behavior not specified by IEEE Std. 1003.1-200x which
221 results from use of a valid program construct or valid data input.

222 The value or behavior may vary among implementations that conform to
223 IEEE Std. 1003.1-200x. An application should not rely on the existence or validity of the
224 value or behavior. An application that relies on any particular value or behavior cannot be
225 assured to be portable across conforming implementations.

226 **1.7 Definitions**

227 Concepts and definitions are defined in the Base Definitions volume of IEEE Std. 1003.1-200x.

228 1.8 Relationship to Other Formal Standards

229 Great care has been taken to ensure that this volume of IEEE Std. 1003.1-200x is fully aligned
230 with the following standards:

231 ISO C (1999)

232 ISO/IEC 9899: 1999, Programming Languages — C.

233 Parts of the ISO/IEC 9899: 1999 standard (hereinafter referred to as the ISO C standard) are
234 referenced to describe requirements also mandated by this volume of IEEE Std. 1003.1-200x.
235 Some functions and headers included within this volume of IEEE Std. 1003.1-200x have a version
236 in the ISO C standard; in this case *CX* markings are added as appropriate to show where the
237 ISO C standard has been extended. Any conflict between this volume of IEEE Std. 1003.1-200x
238 and the ISO C standard is unintentional.

239 This volume of IEEE Std. 1003.1-200x also allows, but does not require, mathematics functions to
240 support IEEE Std. 754-1985 and IEEE Std. 854-1987.

241 1.9 Portability

242 Some of the utilities in the Shell and Utilities volume of IEEE Std. 1003.1-200x and functions in
 243 the System Interfaces volume of IEEE Std. 1003.1-200x describe functionality that might not be
 244 fully portable to systems meeting the requirements for POSIX conformance (see the Base
 245 Definitions volume of IEEE Std. 1003.1-200x, Chapter 2, Conformance).

246 Where optional, enhanced, or reduced functionality is specified, the text is shaded and a code in
 247 the margin identifies the nature of the option, extension, or warning (see Section 1.9.1). For
 248 maximum portability, an application should avoid such functionality.

249 1.9.1 Codes

250 Margin codes and their meanings are listed in the Base Definitions volume of
 251 IEEE Std. 1003.1-200x, but are repeated here for convenience:

252 ADV **Advisory Information**

253 The functionality described is optional. The functionality described is also an extension to the
 254 ISO C standard.

255 Where applicable, functions are marked with the ADV margin legend in the SYNOPSIS section.
 256 Where additional semantics apply to a function, the material is identified by use of the ADV
 257 margin legend.

258 AIO **Asynchronous Input and Output**

259 The functionality described is optional. The functionality described is also an extension to the
 260 ISO C standard.

261 Where applicable, functions are marked with the AIO margin legend in the SYNOPSIS section.
 262 Where additional semantics apply to a function, the material is identified by use of the AIO
 263 margin legend.

264 BAR **Barriers**

265 The functionality described is optional. The functionality described is also an extension to the
 266 ISO C standard.

267 Where applicable, functions are marked with the BAR margin legend in the SYNOPSIS section.
 268 Where additional semantics apply to a function, the material is identified by use of the BAR
 269 margin legend.

270 BE **Batch Environment Services and Utilities**

271 The functionality described is optional.

272 Where applicable, utilities are marked with the BE margin legend in the SYNOPSIS section.
 273 Where additional semantics apply to a utility, the material is identified by use of the BE margin
 274 legend.

275 CD **C-Language Development Utilities**

276 The functionality described is optional.

277 Where applicable, utilities are marked with the CD margin legend in the SYNOPSIS section.
 278 Where additional semantics apply to a utility, the material is identified by use of the CD margin
 279 legend.

280 CPT **Process CPU-Time Clocks**

281 The functionality described is optional. The functionality described is also an extension to the
 282 ISO C standard.

283 Where applicable, functions are marked with the CPT margin legend in the SYNOPSIS section.
 284 Where additional semantics apply to a function, the material is identified by use of the CPT

285 margin legend.

286 CS **Clock Selection**
287 The functionality described is optional. The functionality described is also an extension to the
288 ISO C standard.

289 Where applicable, functions are marked with the CS margin legend in the SYNOPSIS section.
290 Where additional semantics apply to a function, the material is identified by use of the CS
291 margin legend.

292 CX **Extension to the ISO C standard**
293 The functionality described is an extension to the ISO C standard. Application writers may
294 make use of an extension as it is supported on all IEEE Std. 1003.1-200x-conforming systems.

295 FD **FORTTRAN Development Utilities**
296 The functionality described is optional.

297 Where applicable, utilities are marked with the FD margin legend in the SYNOPSIS section.
298 Where additional semantics apply to a utility, the material is identified by use of the FD margin
299 legend.

300 FR **FORTTRAN Runtime Utilities**
301 The functionality described is optional.

302 Where applicable, utilities are marked with the FR margin legend in the SYNOPSIS section.
303 Where additional semantics apply to a utility, the material is identified by use of the FR margin
304 legend.

305 FSC **File Synchronization**
306 The functionality described is optional. The functionality described is also an extension to the
307 ISO C standard.

308 Where applicable, functions are marked with the FSC margin legend in the SYNOPSIS section.
309 Where additional semantics apply to a function, the material is identified by use of the FSC
310 margin legend.

311 IP6 **IPV6**
312 The functionality described is optional. The functionality described is also an extension to the
313 ISO C standard.

314 Where applicable, functions are marked with the IP6 margin legend in the SYNOPSIS section.
315 Where additional semantics apply to a function, the material is identified by use of the IP6
316 margin legend.

317 MAN **Mandatory in the Next Draft**
318 This is an interim draft code used to aid reviewers during the development of
319 IEEE Std. 1003.1-200x. It denotes a feature that was previously an option or extension that is
320 being brought into the mandatory base functionality. This margin code will be removed from the
321 final draft.

322 MF **Memory Mapped Files**
323 The functionality described is optional. The functionality described is also an extension to the
324 ISO C standard.

325 Where applicable, functions are marked with the MF margin legend in the SYNOPSIS section.
326 Where additional semantics apply to a function, the material is identified by use of the MF
327 margin legend.

328 ML **Process Memory Locking**
329 The functionality described is optional. The functionality described is also an extension to the

330 ISO C standard.

331 Where applicable, functions are marked with the ML margin legend in the SYNOPSIS section.
 332 Where additional semantics apply to a function, the material is identified by use of the ML
 333 margin legend.

334 MLR **Range Memory Locking**
 335 The functionality described is optional. The functionality described is also an extension to the
 336 ISO C standard.

337 Where applicable, functions are marked with the MLR margin legend in the SYNOPSIS section.
 338 Where additional semantics apply to a function, the material is identified by use of the MLR
 339 margin legend.

340 MON **Monotonic Clock**
 341 The functionality described is optional. The functionality described is also an extension to the
 342 ISO C standard.

343 Where applicable, functions are marked with the MON margin legend in the SYNOPSIS section.
 344 Where additional semantics apply to a function, the material is identified by use of the MON
 345 margin legend.

346 MPR **Memory Protection**
 347 The functionality described is optional. The functionality described is also an extension to the
 348 ISO C standard.

349 Where applicable, functions are marked with the MPR margin legend in the SYNOPSIS section.
 350 Where additional semantics apply to a function, the material is identified by use of the MPR
 351 margin legend.

352 MSG **Message Passing**
 353 The functionality described is optional. The functionality described is also an extension to the
 354 ISO C standard.

355 Where applicable, functions are marked with the MSG margin legend in the SYNOPSIS section.
 356 Where additional semantics apply to a function, the material is identified by use of the MSG
 357 margin legend.

358 OB **Obsolescent**
 359 The functionality described may be withdrawn in a future version of this volume of
 360 IEEE Std. 1003.1-200x. Strictly Conforming POSIX Applications and Strictly Conforming XSI
 361 Applications shall not use obsolescent features.

362 OF **Output Format Incompletely Specified**
 363 The functionality described is an XSI extension. The format of the output produced by the utility
 364 is not fully specified. It is therefore not possible to post-process this output in a consistent
 365 fashion. Typical problems include unknown length of strings and unspecified field delimiters.

366 OH **Optional Header**
 367 In the SYNOPSIS section of some interfaces in the System Interfaces volume of
 368 IEEE Std. 1003.1-200x an included header is marked as in the following example:

369 OH `#include <sys/types.h>`
 370 `#include <grp.h>`
 371 `struct group *getgrnam(const char *name);`

372 This indicates that the marked header is not required on XSI-conformant systems.

373	PIO	Prioritized Input and Output
374		The functionality described is optional. The functionality described is also an extension to the
375		ISO C standard.
376		Where applicable, functions are marked with the PIO margin legend in the SYNOPSIS section.
377		Where additional semantics apply to a function, the material is identified by use of the PIO
378		margin legend.
379	PS	Process Scheduling
380		The functionality described is optional. The functionality described is also an extension to the
381		ISO C standard.
382		Where applicable, functions are marked with the PS margin legend in the SYNOPSIS section.
383		Where additional semantics apply to a function, the material is identified by use of the PS
384		margin legend.
385	RTS	Realtime Signals Extension
386		The functionality described is optional. The functionality described is also an extension to the
387		ISO C standard.
388		Where applicable, functions are marked with the RTS margin legend in the SYNOPSIS section.
389		Where additional semantics apply to a function, the material is identified by use of the RTS
390		margin legend.
391	SD	Software Development Utilities
392		The functionality described is optional.
393		Where applicable, utilities are marked with the SD margin legend in the SYNOPSIS section.
394		Where additional semantics apply to a utility, the material is identified by use of the SD margin
395		legend.
396	SEM	Semaphores
397		The functionality described is optional. The functionality described is also an extension to the
398		ISO C standard.
399		Where applicable, functions are marked with the SEM margin legend in the SYNOPSIS section.
400		Where additional semantics apply to a function, the material is identified by use of the SEM
401		margin legend.
402	SHM	Shared Memory Objects
403		The functionality described is optional. The functionality described is also an extension to the
404		ISO C standard.
405		Where applicable, functions are marked with the SHM margin legend in the SYNOPSIS section.
406		Where additional semantics apply to a function, the material is identified by use of the SHM
407		margin legend.
408	SIO	Synchronized Input and Output
409		The functionality described is optional. The functionality described is also an extension to the
410		ISO C standard.
411		Where applicable, functions are marked with the SIO margin legend in the SYNOPSIS section.
412		Where additional semantics apply to a function, the material is identified by use of the SIO
413		margin legend.
414	SPI	Spin Locks
415		The functionality described is optional. The functionality described is also an extension to the
416		ISO C standard.

417 Where applicable, functions are marked with the SPI margin legend in the SYNOPSIS section.
418 Where additional semantics apply to a function, the material is identified by use of the SPI
419 margin legend.

420 SPN **Spawn**

421 The functionality described is optional. The functionality described is also an extension to the
422 ISO C standard.

423 Where applicable, functions are marked with the SPN margin legend in the SYNOPSIS section.
424 Where additional semantics apply to a function, the material is identified by use of the SPN
425 margin legend.

426 SS **Process Sporadic Server**

427 The functionality described is optional. The functionality described is also an extension to the
428 ISO C standard.

429 Where applicable, functions are marked with the SS margin legend in the SYNOPSIS section.
430 Where additional semantics apply to a function, the material is identified by use of the SS
431 margin legend.

432 TCT **Thread CPU-Time Clocks**

433 The functionality described is optional. The functionality described is also an extension to the
434 ISO C standard.

435 Where applicable, functions are marked with the TCT margin legend in the SYNOPSIS section.
436 Where additional semantics apply to a function, the material is identified by use of the TCT
437 margin legend.

438 THR **Threads**

439 The functionality described is optional. The functionality described is also an extension to the
440 ISO C standard.

441 Where applicable, functions are marked with the THR margin legend in the SYNOPSIS section.
442 Where additional semantics apply to a function, the material is identified by use of the THR
443 margin legend.

444 TMO **Timeouts**

445 The functionality described is optional. The functionality described is also an extension to the
446 ISO C standard.

447 Where applicable, functions are marked with the TMO margin legend in the SYNOPSIS section.
448 Where additional semantics apply to a function, the material is identified by use of the TMO
449 margin legend.

450 TMR **Timers**

451 The functionality described is optional. The functionality described is also an extension to the
452 ISO C standard.

453 Where applicable, functions are marked with the TMR margin legend in the SYNOPSIS section.
454 Where additional semantics apply to a function, the material is identified by use of the TMR
455 margin legend.

456 TPI **Threads Priority Inheritance**

457 The functionality described is optional. The functionality described is also an extension to the
458 ISO C standard.

459 Where applicable, functions are marked with the TPI margin legend in the SYNOPSIS section.
460 Where additional semantics apply to a function, the material is identified by use of the TPI
461 margin legend.

462 TPP **Thread Priority Protection**
463 The functionality described is optional. The functionality described is also an extension to the
464 ISO C standard.

465 Where applicable, functions are marked with the TPP margin legend in the SYNOPSIS section.
466 Where additional semantics apply to a function, the material is identified by use of the TPP
467 margin legend.

468 TPS **Thread Execution Scheduling**
469 The functionality described is optional. The functionality described is also an extension to the
470 ISO C standard.

471 Where applicable, functions are marked with the TPS margin legend for the SYNOPSIS section.
472 Where additional semantics apply to a function, the material is identified by use of the TPS
473 margin legend.

474 TRC **Trace**
475 The functionality described is optional. The functionality described is also an extension to the
476 ISO C standard.

477 Where applicable, functions are marked with the TRC margin legend in the SYNOPSIS section.
478 Where additional semantics apply to a function, the material is identified by use of the TRC
479 margin legend.

480 TEF **Trace Event Filter**
481 The functionality described is optional. The functionality described is also an extension to the
482 ISO C standard.

483 Where applicable, functions are marked with the TEF margin legend in the SYNOPSIS section.
484 Where additional semantics apply to a function, the material is identified by use of the TEF
485 margin legend.

486 TRL **Trace Log**
487 The functionality described is optional. The functionality described is also an extension to the
488 ISO C standard.

489 Where applicable, functions are marked with the TRL margin legend in the SYNOPSIS section.
490 Where additional semantics apply to a function, the material is identified by use of the TRL
491 margin legend.

492 TRI **Trace Inherit**
493 The functionality described is optional. The functionality described is also an extension to the
494 ISO C standard.

495 Where applicable, functions are marked with the TRI margin legend in the SYNOPSIS section.
496 Where additional semantics apply to a function, the material is identified by use of the TRI
497 margin legend.

498 TSA **Thread Stack Address Attribute**
499 The functionality described is optional. The functionality described is also an extension to the
500 ISO C standard.

501 Where applicable, functions are marked with the TPS margin legend for the SYNOPSIS section.
502 Where additional semantics apply to a function, the material is identified by use of the TSA
503 margin legend.

504 TSF **Thread-Safe Functions**
505 The functionality described is optional. The functionality described is also an extension to the
506 ISO C standard.

507 Where applicable, functions are marked with the TSF margin legend in the SYNOPSIS section.
508 Where additional semantics apply to a function, the material is identified by use of the TSF
509 margin legend.

510 TSH **Thread Process-Shared Synchronization**
511 The functionality described is optional. The functionality described is also an extension to the
512 ISO C standard.

513 Where applicable, functions are marked with the TSH margin legend in the SYNOPSIS section.
514 Where additional semantics apply to a function, the material is identified by use of the TSH
515 margin legend.

516 TSP **Thread Sporadic Server**
517 The functionality described is optional. The functionality described is also an extension to the
518 ISO C standard.

519 Where applicable, functions are marked with the TSP margin legend in the SYNOPSIS section.
520 Where additional semantics apply to a function, the material is identified by use of the TSP
521 margin legend.

522 TSS **Thread Stack Address Size**
523 The functionality described is optional. The functionality described is also an extension to the
524 ISO C standard.

525 Where applicable, functions are marked with the TSS margin legend in the SYNOPSIS section.
526 Where additional semantics apply to a function, the material is identified by use of the TSS
527 margin legend.

528 TYM **Typed Memory Objects**
529 The functionality described is optional. The functionality described is also an extension to the
530 ISO C standard.

531 Where applicable, functions are marked with the TYM margin legend in the SYNOPSIS section.
532 Where additional semantics apply to a function, the material is identified by use of the TYM
533 margin legend.

534 UN **Possibly Unsupportable Feature**
535 The functionality described is an XSI extension. It need not be possible to implement the
536 required functionality (as defined) on all conformant systems and the functionality need not be
537 present. This may, for example, be the case where the conformant system is hosted and the
538 underlying system provides the service in an alternative way.

539 UP **User Portability Utilities**
540 The functionality described is optional.

541 Where applicable, utilities are marked with the UP margin legend in the SYNOPSIS section.
542 Where additional semantics apply to a utility, the material is identified by use of the UP margin
543 legend.

544 XSI **Extension**
545 The functionality described is an XSI extension. Functionality marked XSI is also an extension to
546 the ISO C standard. Application writers may confidently make use of an extension on all
547 systems supporting the X/Open System Interfaces Extension.

548 If an entire SYNOPSIS section is shaded and marked with one XSI, all the functionality described
549 in that reference page is an extension. See the Base Definitions volume of IEEE Std. 1003.1-200x,
550 Section 3.441, XSI.

551 XSR **XSI STREAMS**
552 The functionality described is optional. The functionality described is also an extension to the
553 ISO C standard.

554 Where applicable, functions are marked with the XSR margin legend in the SYNOPSIS section.
555 Where additional semantics apply to a function, the material is identified by use of the XSR
556 margin legend.

557 1.10 Format of Entries

558 The entries in Chapter 3 are based on a common format as follows. The only sections relating to
559 conformance are the SYNOPSIS, DESCRIPTION, RETURN VALUE, and ERRORS sections.

560 NAME

561 This section gives the name or names of the entry and briefly states its purpose.

562 SYNOPSIS

563 This section summarizes the use of the entry being described. If it is necessary to
564 include a header to use this function, the names of such headers are shown, for
565 example:

```
566 #include <stdio.h>
```

567 DESCRIPTION

568 This section describes the functionality of the function or header.

569 RETURN VALUE

570 This section indicates the possible return values, if any.

571 If the implementation can detect errors, “successful completion” means that no error
572 has been detected during execution of the function. If the implementation does detect
573 an error, the error is indicated.

574 For functions where no errors are defined, “successful completion” means that if the
575 implementation checks for errors, no error has been detected. If the implementation can
576 detect errors, and an error is detected, the indicated return value is returned and *errno*
577 may be set.

578 ERRORS

579 This section gives the symbolic names of the values returned in *errno* if an error occurs.

580 “No errors are defined” means that values and usage of *errno*, if any, depend on the
581 implementation.

582 EXAMPLES

583 This section is non-normative.

584 This section gives examples of usage, where appropriate. In the event of conflict
585 between an example and a normative part of this volume of IEEE Std. 1003.1-200x, the
586 normative material is to be taken as correct.

587 APPLICATION USAGE

588 This section is non-normative.

589 This section gives warnings and advice to application writers about the entry. In the
590 event of conflict between warnings and advice and a normative part of this volume of
591 IEEE Std. 1003.1-200x, the normative material is to be taken as correct.

592 RATIONALE

593 This section is non-normative.

594 This section contains historical information concerning the contents of this volume of
595 IEEE Std. 1003.1-200x and why features were included or discarded by the standard
596 developers.

597 FUTURE DIRECTIONS

598 This section is non-normative.

599 This section provides comments which should be used as a guide to current thinking;
600 there is not necessarily a commitment to adopt these future directions.

601 **SEE ALSO**

602 This section is non-normative.

603 This section gives references to related information.

604 **CHANGE HISTORY**

605 This section is non-normative.

606 This section shows the derivation of the entry and any significant changes that have
607 been made to it.

608

609 This chapter covers information that is relevant to all the functions specified in Chapter 3 and
610 the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 13, Headers.

611 **2.1 Use and Implementation of Functions**

612 Each of the following statements shall apply unless explicitly stated otherwise in the detailed
613 descriptions that follow:

- 614 1. If an argument to a function has an invalid value (such as a value outside the domain of
615 the function, or a pointer outside the address space of the program, or a null pointer), the
616 behavior is undefined.
- 617 2. Any function declared in a header may also be implemented as a macro defined in the
618 header, so a library function should not be declared explicitly if its header is included. Any
619 macro definition of a function can be suppressed locally by enclosing the name of the
620 function in parentheses, because the name is then not followed by the left parenthesis that
621 indicates expansion of a macro function name. For the same syntactic reason, it is
622 permitted to take the address of a library function even if it is also defined as a macro. The
623 use of the C-language **#undef** construct to remove any such macro definition shall also
624 ensure that an actual function is referred to.
- 625 3. Any invocation of a library function that is implemented as a macro shall expand to code
626 that evaluates each of its arguments exactly once, fully protected by parentheses where
627 necessary, so it is generally safe to use arbitrary expressions as arguments. Likewise, those
628 function-like macros described in the following sections may be invoked in an expression
629 anywhere a function with a compatible return type could be called.
- 630 4. Provided that a library function can be declared without reference to any type defined in a
631 header, it is also permissible to declare the function, either explicitly or implicitly, and use
632 it without including its associated header.
- 633 5. If a function that accepts a variable number of arguments is not declared (explicitly or by
634 including its associated header), the behavior is undefined.

635 2.2 The Compilation Environment

636 2.2.1 POSIX.1 Symbols

637 Certain symbols in this volume of IEEE Std. 1003.1-200x are defined in headers (see the Base
638 Definitions volume of IEEE Std. 1003.1-200x, Chapter 13, Headers). Some of those headers could
639 also define other symbols than those defined by this volume of IEEE Std. 1003.1-200x, potentially
640 conflicting with symbols used by the application. Also, this volume of IEEE Std. 1003.1-200x
641 defines symbols that are not permitted by other standards to appear in those headers without
642 some control on the visibility of those symbols.

643 Symbols called feature test macros are used to control the visibility of symbols that might be
644 included in a header. Implementations, future versions of this volume of IEEE Std. 1003.1-200x,
645 and other standards may define additional feature test macros.

646 In the compilation of an application that **#defines** a feature test macro specified by
647 IEEE Std. 1003.1-200x, no header defined by IEEE Std. 1003.1-200x shall be included prior to the
648 definition of the feature test macro. This restriction also applies to any implementation-
649 provided header in which these feature test macros are used. If the definition of the macro does
650 not precede the **#include**, the result is undefined.

651 Feature test macros shall begin with the underscore character ('_').

652 2.2.1.1 The `_POSIX_C_SOURCE` Feature Test Macro

653 A POSIX-conforming application should ensure that the feature test macro `_POSIX_C_SOURCE`
654 is defined before inclusion of any header.

655 When an application includes a header described by this volume of IEEE Std. 1003.1-200x, and
656 when this feature test macro is defined to have at least the value `200xMML`:

- 657 1. All symbols required by this volume of IEEE Std. 1003.1-200x to appear when the header is
658 included shall be made visible.
- 659 2. Symbols that are explicitly permitted, but not required, by this volume of
660 IEEE Std. 1003.1-200x to appear in that header (including those in reserved name spaces)
661 may be made visible.
- 662 3. Additional symbols not required or explicitly permitted by this volume of
663 IEEE Std. 1003.1-200x to be in that header shall not be made visible, except when enabled
664 by another feature test macro or by having defined `_POSIX_C_SOURCE` with a value
665 larger than `200xxxL`.

666 Identifiers in this volume of IEEE Std. 1003.1-200x may only be undefined using the **#undef**
667 directive as described in Section 2.1 (on page 511) or Section 2.2.2 (on page 513). These **#undef**
668 directives shall follow all **#include** directives of any header in this volume of
669 IEEE Std. 1003.1-200x.

670 2.2.1.2 The `_XOPEN_SOURCE` Feature Test Macro

671 XSI An XSI-conforming application should ensure that the feature test macro `_XOPEN_SOURCE` is
672 defined with the value `600` before inclusion of any header. This is needed to enable the
673 functionality described in Section 2.2.1.1 and in addition to enable the X/Open System Interfaces
674 Extension.

675 Since this volume of IEEE Std. 1003.1-200x is aligned with the ISO C standard, and since all
676 functionality enabled by `_POSIX_C_SOURCE` set greater than zero and less than or equal to
677 `200xxxL` should be enabled by `_XOPEN_SOURCE` set equal to `600`, there should be no need to

678 define either `_POSIX_SOURCE` or `_POSIX_C_SOURCE` if `_XOPEN_SOURCE` is so defined.
679 Therefore, if `_XOPEN_SOURCE` is set equal to 600 and `_POSIX_SOURCE` is defined, or
680 `_POSIX_C_SOURCE` is set greater than zero and less than or equal to 200xxxL, the behavior is
681 the same as if only `_XOPEN_SOURCE` is defined and set equal to 600. However, should
682 `_POSIX_C_SOURCE` be set to a value greater than 200xxxL, the behavior is undefined.

683 2.2.2 The Name Space

684 All identifiers in this volume of IEEE Std. 1003.1-200x, except *environ*, are defined in at least one
685 of the headers, as shown in the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 13,
686 XSI Headers. When `_XOPEN_SOURCE` or `_POSIX_C_SOURCE` is defined, each header defines or
687 declares some identifiers, potentially conflicting with identifiers used by the application. The set
688 of identifiers visible to the application consists of precisely those identifiers from the header
689 pages of the included headers, as well as additional identifiers reserved for the implementation.
690 In addition, some headers may make visible identifiers from other headers as indicated on the
691 relevant header pages.

692 Implementations may also add members to a structure or union without controlling the
693 visibility of those members with a feature test macro, as long as a user-defined macro with the
694 same name cannot interfere with the correct interpretation of the program. The identifiers
695 reserved for use by the implementation are described below:

- 696 1. Each identifier with external linkage described in the header section is reserved for use as
697 an identifier with external linkage if the header is included.
- 698 2. Each macro name described in the header section is reserved for any use if the header is
699 included.
- 700 3. Each identifier with file scope described in the header section is reserved for use as an
701 identifier with file scope in the same name space if the header is included.

702 The prefixes `posix_`, `POSIX_`, and `_POSIX_` are reserved for use by IEEE Std. 1003.1-200x and
703 other POSIX standards. Implementations may add symbols to the headers shown in the
704 following table, provided the identifiers for those symbols begin with the corresponding
705 reserved prefixes in the following table, and do not use the reserved prefixes `posix_`, `POSIX_`, or
706 `_POSIX_`.

707
708
709
710 AIO
711
712
713
714
715
716
717
718
719
720
721 MSG
722 XSI
723
724
725
726 XSI
727
728
729
730 PS
731 SEM
732
733 XSI
734 RTS
735 XSI
736
737
738
739 XSI
740 MF
741 XSI
742 XSI
743 XSI
744 XSI
745
746
747
748 XSI
749
750
751 XSI
752
753
754 XSI
755

Header	Prefix	Suffix	Complete Name
<aio.h>	aio_, lio_, AIO_, LIO_		
<arpa/inet.h>	in_, inet_		
<ctype.h>	to[a-z], is[a-z]		
<dirent.h>	d_		
<errno.h>	E[0-9], E[A-Z]		
<fcntl.h>	l_		
<glob.h>	gl_		
<grp.h>	gr_		
<inttypes.h>	int[0-9a-z_]_t, uint[0-9a-z_]_t		
<limits.h>		_MAX, _MIN	
<locale.h>	LC_[A-Z]		
<mqueue.h>	mq_, MQ_		
<ndbm.h>	dbm_		
<netdb.h>	h_, n_, p_, s_		
<net/if.h>	if_		
<netinet/in.h>	in_, ip_, s_, sin_		
<poll.h>	pd_, ph_, ps_		
<pthread.h>	pthread_, PTHREAD_		
<pwd.h>	pw_		
<regex.h>	re_, rm_		
<sched.h>	sched_, SCHED_		
<semaphore.h>	sem_, SEM_		
<signal.h>	sa_, uc_, SIG[A-Z], SIG_[A-Z]		
	ss_, sv_		
	si_, SI_, sigev_, SIGEV_, sival_		
<stropts.h>	bi_, ic_, l_, sl_, str_		
<stdint.h>	int[0-9a-z_]_t, uint[0-9a-z_]_t		
<stdlib.h>	str[a-z]		
<string.h>	str[a-z], mem[a-z], wcs[a-z]		
<sys/ipc.h>	ipc_		key, pad, seq
<sys/mman.h>	shm_, MAP_, MCL_, MS_, PROT_		
<sys/msg.h>	msg		msg
<sys/resource.h>	rlim_, ru_		
<sys/sem.h>	sem		sem
<sys/shm.h>	shm		
<sys/socket.h>	_ss, sa_, if_, ifc_, ifru_, infu_, ifra_, msg_, cmsg_, l_		
<sys/stat.h>	st_		
<sys/statvfs.h>	f_		
<sys/time.h>	fds_, it_, tv_, FD_		
<sys/times.h>	tms_		
<sys/uio.h>	iov_		
<sys/un.h>	sun_		
<sys/utsname.h>	uts_		
<sys/wait.h>	si_, W[A-Z], P_		
<termios.h>	c_		

756
 757
 758
 759
 760 TMR
 761 TMR
 762 XSI
 763 XSI
 764
 765 XSI
 766
 767
 768
 769
 770

Header	Prefix	Suffix	Complete Name
<time.h>	tm_		
	clock_, timer_, it_, tv_, CLOCK_, TIMER_		
<ucontext.h>	uc_, ss_		
<ulimit.h>	UL_		
<utime.h>	utim_		
<utmpx.h>	ut_	_LVL, _TIME, _PROCESS	
<wchar.h>	wcs[a-z]		
<wctype.h>	is[a-z], to[a-z]		
<wordexp.h>	we_		
ANY header	POSIX_, _POSIX_, posix_	_t	

771 **Note:** The notation [A-Z] indicates any uppercase letter in the portable character set. The
 772 notation [a-z] indicates any lowercase letter in the portable character set. Commas
 773 and spaces in the lists of prefixes and complete names in the above table are not part
 774 of any prefix or complete name.

775 If any header in the following table is included, macros with the prefixes shown may be defined. |
 776 After the last inclusion of a given header, an application may use identifiers with the
 777 corresponding prefixes for its own purpose, provided their use is preceded by an **#undef** of the
 778 corresponding macro.

779
780
781 XSI
782
783 XSI
784
785 XSI
786
787
788 XSI
789
790
791
792 XSI
793 XSI
794
795
796 XSI
797
798
799 XSI
800 XSI
801 XSI
802 XSI
803 XSI
804 XSI
805 XSI
806 XSI
807 XSI
808
809 XSI
810 XSI
811 XSI
812 XSI
813
814

Header	Prefix
<dfcn.h>	RTLD_
<fcntl.h>	F_, O_, S_
<fmtmsg.h>	MM_
<fnmatch.h>	FNM_
<ftw.h>	FTW
<glob.h>	GLOB_
<inttypes.h>	PRI[a-z], SCN[a-z]
<ndbm.h>	DBM_
<net/if.h>	IF_
<netinet/in.h>	IMPLINK_, IN_, INADDR_, IP_, IPPORT_, IPPROTO_, SOCK_
<netinet/tcp.h>	TCP_
<nl_types.h>	NL_
<poll.h>	POLL
<regex.h>	REG_
<signal.h>	SA_, SIG_[0-9a-z_],
	BUS_, CLD_, FPE_, ILL_, POLL_, SEGV_, SI_, SS_, SV_, TRAP_
stdint.h	INT[0-9A-Z_]_MIN, INT[0-9A-Z_]_MAX, INT[0-9A-Z_]_C
	UINT[0-9A-Z_]_MIN, UINT[0-9A-Z_]_MAX, UINT[0-9A-Z_]_C
<stropts.h>	FLUSH[A-Z], I_, M_, MUXID_R[A-Z], S_, SND[A-Z], STR
<syslog.h>	LOG_
<sys/ipc.h>	IPC_
<sys/mman.h>	PROT_, MAP_, MS_
<sys/msg.h>	MSG[A-Z]
<sys/resource.h>	PRIO_, RLIM_, RLIMIT_, RUSAGE_
<sys/sem.h>	SEM_
<sys/shm.h>	SHM[A-Z], SHM_[A-Z]
<sys/socket.h>	AF_, CMSG_, MSG_, PF_, SCM_, SHUT_, SO
<sys/stat.h>	S_
<sys/statvfs.h>	ST_
<sys/time.h>	FD_, ITIMER_
<sys/uio.h>	IOV_
<sys/wait.h>	BUS_, CLD_, FPE_, ILL_, POLL_, SEGV_, SI_, TRAP_
<termios.h>	V, I, O, TC, B[0-9] (See below.)
<wordexp.h>	WRDE_

815 **Note:** The notation [0–9] indicates any digit. The notation [A–Z] indicates any uppercase
816 letter in the portable character set. The notation [0–9a–z_] indicates any digit, any
817 lowercase letter in the portable character set, or underscore.

818 The following reserved names are used as exact matches for <termios.h>:

819 XSI	CBAUD	EXTB	VDSUSP
820	DEFECHO	FLUSHO	VLNEXT
821	ECHOCTL	LOBLK	VREPRINT
822	ECHOKE	PENDIN	VSTATUS
823	ECHOPRT	SWTCH	VWERASE
824	EXTA	VDISCARD	

- 825 The following identifiers are reserved regardless of the inclusion of headers:
- 826 1. All identifiers that begin with an underscore and either an uppercase letter or another
827 underscore are always reserved for any use by the implementation.
- 828 2. All identifiers that begin with an underscore are always reserved for use as identifiers with
829 file scope in both the ordinary identifier and tag name spaces.
- 830 3. All identifiers in the table below are reserved for use as identifiers with external linkage.
831 Some of these identifiers do not appear in this volume of IEEE Std. 1003.1-200x, but are
832 reserved for future use by the ISO C standard.

833	_Exit	chrtf	exit	fsetpos	mbrtowc	sinhl
834	abort	chrtl	exp	ftell	mbsinit	sinl
835	abs	cqos	exp2	fwide	mbsrtowcs	sprintf
836	acos	cqosf	exp2f	fwprintf	mbstowcs	sqrt
837	acosf	cqosh	exp2l	fwwrite	mbtowc	sqrtf
838	acosh	cqoshf	expf	fwscanf	memb[a-z]*	sqrtl
839	acoshf	cqoshl	expl	getc	mktime	srand
840	acoshl	cqosl	expm1	getchar	modf	sscanf
841	acosl	ceil	expm1f	getenv	modff	str[a-z]*
842	acosl	ceilf	expm1l	gets	modfl	strtof
843	asctime	ceilf	fabs	getwc	nan	strtoimax
844	asin	ceil	fabsf	getwchar	nanf	strtold
845	asinf	ceil	fabsl	gmtime	nanl	strtoll
846	asinh	cexp	fclose	hypotf	nearbyint	strtoull
847	asinhf	cexpf	fdim	hypotl	nearbyintf	strtoumax
848	asinhl	cexpl	fdimf	ilogb	nearbyintl	swprintf
849	asinl	cimag	fdiml	ilogbf	nextafterf	swscanf
850	asinl	cimagf	fclearexcept	ilogbl	nextafterl	system
851	atan	cimagl	fgetenv	imaxabs	nexttoward	tan
852	atan2	clearerr	fgetexceptflag	imaxdiv	nexttowardf	tanf
853	atan2f	clock	fgetround	is[a-z]*	nexttowardl	tanh
854	atan2l	clog	fholdexcept	isblank	perror	tanhf
855	atanf	clogf	fEOF	iswblank	pow	tanhl
856	atanf	clogl	fraiseexcept	labs	powf	tanl
857	atanh	conj	ferror	ldexp	powl	tgamma
858	atanh	conjf	fsetenv	ldexpf	printf	tgammaf
859	atanhf	conjl	fsetexceptflag	ldexpl	putc	tgammal
860	atanhl	copysign	fsetround	ldiv	putchar	time
861	atanl	copysignf	fetestexcept	ldiv	puts	tmpfile
862	atanl	copysignl	fupdateenv	lgammaf	putwc	tmpnam
863	atexit	cqs	fflush	lgammal	putwchar	to[a-z]*
864	atof	cqsf	fgetc	llabs	qsort	trunc
865	atoi	cqsh	fgetpos	llrint	raise	truncf
866	atol	cqshf	fgets	llrintf	rand	truncl
867	atoll	cqshl	fgetwc	llrintl	realloc	ungetc
868	bsearch	cqsl	fgetws	llround	remainderf	ungetwc
869	cabs	cpow	floor	llroundf	remainderl	va_end
870	cabsf	cpowf	floorf	llroundl	remove	vfprintf
871	cabsl	cpowl	floorl	lqcaleconv	remquo	vscanf
872	cacos	cproj	fma	lqcaltime	remquof	vfprintf

Notes to Reviewers

This section with side shading will not appear in the final copy. - Ed.

The following table should be made complete by including everything not in the previous table.

- The following identifiers are also reserved for use as identifiers with external linkage:

Table 2-1 XSI Identifiers

a64l	fcntl	getpriority	mknod	regex	sigelse
basename	fcntl	getpwent	mkstemp	remainder	sigset
bcmp	fchmod	getrlimit	mktemp	remque	sigstack
bcopy	fchown	getrusage	mmap	rindex	srandom
brk	fcntl	getsid	mprotect	sbrk	statvfs
bsd_signal	fcntl	getsubopt	rand48	scalb	strcascmp
bzero	ffs	gettimeofday	msync	select	strdup
cbrt	fcntl	getutxent	mmap	setcontext	strcascmp
closelog	statvfs	getutxid	nextafter	setgrent	swapcontext
dbm_clearerr	ftime	getutxline	nftw	setitimer	symlink
dbm_close	ftok	getwd	nicp	_setjmp	sync
dbm_delete	truncate	grantpt	openlog	setlogmask	syslog
dbm_error	gcvt	index	poll	setpgpr	tcgetsid
dbm_fetch	getcontext	initsstate	ptname	setpriority	truncate
dbm_firstkey	getdate	insque	putmsg	setpwent	ttyslot
dbm_nextkey	getdtablesize	ioctl	putpmsg	setreuid	ualarm
dbm_open	getgrent	isastream	pututxline	setrlimit	unlockpt
dbm_store	getgrgid	killpg	random	setstate	usleep
dirname	gethostid	l64a	readlink	setutxent	utimes
ecvt	getitimer	lchown	readv	sigaltstack	valloc
endgrent	getmsg	lockf	realpath	sighold	vfork
endpwent	getpagesize	_longjmp	re_comp	sigignore	wait3
endservent	getpgid	lstat	re_exec	siginterrupt	waitid
endutxent	getpmsg	makecontext	regcmp	sigpause	wrtv

Table 2-2 Sockets Identifiers

accept	if_freenameindex	recvfrom	shutdown
bind	if_indexname	recvmsg	socket
connect	if_nameindex	send	socketpair
getpeername	if_nameindex	sendmsg	
getsockname	listen	sendto	
getsockopt	recv	setsockopt	

938

Table 2-3 IP Address Resolution Identifiers

939	endhostent	getipnodebyaddr	getservbyname	inet_netof
940	endnetent	getipnodebyname	getservbyport	inet_network
941	endprotoent	getnameinfo	getservent	inet_ntoa
942	endservent	getnetbyaddr	h_errno	ntohl
943	getaddrinfo	getnetbyname	htonl	ntohs
944	gethostbyaddr	getnetent	htons	sethostent
945	gethostbyname	getprotobyname	inet_addr	setnetent
946	gethostent	getprotobynumber	inet_lnaof	setprotoent
947	gethostname	getprotoent	inet_makeaddr	setservent

948 All the identifiers defined in this volume of IEEE Std. 1003.1-200x that have external linkage are
 949 always reserved for use as identifiers with external linkage.

950 No other identifiers are reserved.

951 Applications shall not declare or define identifiers with the same name as an identifier reserved
 952 in the same context. Since macro names are replaced whenever found, independent of scope and
 953 name space, macro names matching any of the reserved identifier names shall not be defined by
 954 an application if any associated header is included.

955 Except that the effect of each inclusion of `<assert.h>` depends on the definition of `NDEBUG`,
 956 headers may be included in any order, and each may be included more than once in a given
 957 scope, with no difference in effect from that of being included only once.

958 If used, the application shall ensure that a header is included outside of any external declaration
 959 or definition, and it shall be first included before the first reference to any type or macro it
 960 defines, or to any function or object it declares. However, if an identifier is declared or defined in
 961 more than one header, the second and subsequent associated headers may be included after the
 962 initial reference to the identifier. Prior to the inclusion of a header, the application shall not
 963 define any macros with names lexically identical to symbols defined by that header.

964 2.3 Error Numbers

965 Most functions can provide an error number. The means by which each function provides its
966 error numbers is specified in its description.

967 Some functions provide the error number in a variable accessed through the symbol *errno*. The
968 symbol *errno*, defined by including the `<errno.h>` header, is a macro that expands to a
969 modifiable **lvalue** of type **int**.

970 The value of *errno* should only be examined when it is indicated to be valid by a function's return
971 value. No function in this volume of IEEE Std. 1003.1-200x shall set *errno* to zero. For each thread
972 of a process, the value of *errno* shall not be affected by function calls or assignments to *errno* by
973 other threads.

974 Some functions return an error number directly as the function value. These functions return a
975 value of zero to indicate success.

976 If more than one error occurs in processing a function call, any one of the possible errors may be
977 returned, as the order of detection is undefined.

978 Implementations may support additional errors not included in this list, may generate errors
979 included in this list under circumstances other than those described here, or may contain
980 extensions or limitations that shall prevent some errors from occurring. The ERRORS section on
981 each page specifies whether an error shall be returned, or whether it may be returned.
982 Implementations shall not generate a different error number from the ones described here for
983 error conditions described in this volume of IEEE Std. 1003.1-200x, but may generate additional
984 errors unless explicitly disallowed for a particular function.

985 Each implementation shall document, in the conformance document, situations in which each of
986 the optional conditions defined in IEEE Std. 1003.1-200x are detected. The conformance
987 document may also contain statements that one or more of the optional error conditions are not
988 detected.

989 For functions under the Threads option for which [EINTR] is not listed as a possible error
990 condition in this volume of IEEE Std. 1003.1-200x, an implementation shall not return an error
991 code of [EINTR].

992 The following symbolic names identify the possible error numbers, in the context of the
993 functions specifically defined in this volume of IEEE Std. 1003.1-200x; these general descriptions
994 are more precisely defined in the ERRORS sections of the functions that return them. Only these
995 symbolic names should be used in programs, since the actual value of the error number is
996 unspecified. All values listed in this section shall be unique integer constant expressions with
997 type **int** suitable for use in **#if** preprocessing directives, except as noted below. The values for all
998 these names shall be found in the `<errno.h>` header defined in the Base Definitions volume of
999 IEEE Std. 1003.1-200x. The actual values are unspecified by this volume of IEEE Std. 1003.1-200x.

1000 [E2BIG]

1001 Argument list too long. The sum of the number of bytes used by the new process image's
1002 argument list and environment list is greater than the system-imposed limit of {ARG_MAX}
1003 bytes.

1004 or:

1005 Lack of space in an output buffer.

1006 or:

1007 Argument is greater than the system-imposed maximum.

1008	[EACCES]	
1009		Permission denied. An attempt was made to access a file in a way forbidden by its file
1010		access permissions.
1011	[EADDRINUSE]	
1012		Address in use. The specified address is in use.
1013	[EADDRNOTAVAIL]	
1014		Address not available. The specified address is not available from the local system.
1015	[EAFNOSUPPORT]	
1016		Address family not supported. The implementation does not support the specified address
1017		family, or the specified address is not a valid address for the address family of the specified
1018		socket.
1019	[EAGAIN]	
1020		Resource temporarily unavailable. This is a temporary condition and later calls to the same
1021		routine may complete normally.
1022	[EALREADY]	
1023		Connection already in progress. A connection request is already in progress for the specified
1024		socket.
1025	[EBADF]	
1026		Bad file descriptor. A file descriptor argument is out of range, refers to no open file, or a
1027		read (write) request is made to a file that is only open for writing (reading).
1028	[EBADMSG]	
1029	XSR	Bad message. During a <i>read()</i> , <i>getmsg()</i> , or <i>ioctl()</i> I_RECVFD request to a STREAMS device,
1030		a message arrived at the head of the STREAM that is inappropriate for the function
1031		receiving the message.
1032		<i>read()</i> Message waiting to be read on a STREAM is not a data message.
1033		<i>getmsg()</i> A file descriptor was received instead of a control message.
1034		<i>ioctl()</i> Control or data information was received instead of a file descriptor when
1035		I_RECVFD was specified.
1036		or:
1037		Bad Message. The implementation has detected a corrupted message.
1038	[EBUSY]	
1039		Resource busy. An attempt was made to make use of a system resource that is not currently
1040		available, as it is being used by another process in a manner that would have conflicted with
1041		the request being made by this process.
1042	[ECANCELED]	
1043		Operation canceled. The associated asynchronous operation was canceled before
1044		completion.
1045	[ECHILD]	
1046		No child process. A <i>wait()</i> or <i>waitpid()</i> function was executed by a process that had no
1047		existing or unwaited-for child process.
1048	[ECONNABORTED]	
1049		Connection aborted. The connection has been aborted.
1050	[ECONNREFUSED]	
1051		Connection refused. An attempt to connect to a socket was refused because there was no

1052	process listening or because the queue of connection requests was full and the underlying
1053	protocol does not support retransmissions.
1054	[ECONNRESET]
1055	Connection reset. The connection was forcibly closed by the peer.
1056	[EDEADLK]
1057	Resource deadlock would occur. An attempt was made to lock a system resource that
1058	would have resulted in a deadlock situation.
1059	[EDESTADDRREQ]
1060	Destination address required. No bind address was established.
1061	[EDOM]
1062	Domain error. An input argument is outside the defined domain of the mathematical
1063	function (defined in the ISO C standard).
1064	[EDQUOT]
1065	Reserved.
1066	[EEXIST]
1067	File exists. An existing file was mentioned in an inappropriate context; for example, as a
1068	new link name in the <i>link()</i> function.
1069	[EFAULT]
1070	Bad address. The system detected an invalid address in attempting to use an argument of a
1071	call. The reliable detection of this error cannot be guaranteed, and when not detected may
1072	result in the generation of a signal, indicating an address violation, which is sent to the
1073	process.
1074	[EFBIG]
1075	File too large. The size of a file would exceed the maximum file size of an implementation or
1076	offset maximum established in the corresponding file description.
1077	[EHOSTUNREACH]
1078	Host is unreachable. The destination host cannot be reached (probably because the host is
1079	down or a remote router cannot reach it).
1080	[EIDRM]
1081	Identifier removed. Returned during XSI interprocess communication if an identifier has
1082	been removed from the system.
1083	[EILSEQ]
1084	Illegal byte sequence. A wide-character code has been detected that does not correspond to
1085	a valid character, or a byte sequence does not form a valid wide-character code (defined in
1086	the ISO C standard).
1087	[EINPROGRESS]
1088	Operation in progress. This code is used to indicate that an asynchronous operation has not
1089	yet completed.
1090	or:
1091	O_NONBLOCK is set for the socket file descriptor and the connection cannot be
1092	immediately established.
1093	[EINTR]
1094	Interrupted function call. An asynchronous signal was caught by the process during the
1095	execution of an interruptible function. If the signal handler performs a normal return, the
1096	interrupted function call may return this condition (see the Base Definitions volume of

1097	IEEE Std. 1003.1-200x, <signal.h>).
1098	[EINVAL]
1099	Invalid argument. Some invalid argument was supplied; for example, specifying an
1100	undefined signal in a <i>signal()</i> function or a <i>kill()</i> function.
1101	[EIO]
1102	Input/output error. Some physical input or output error has occurred. This error may be
1103	reported on a subsequent operation on the same file descriptor. Any other error-causing
1104	operation on the same file descriptor may cause the [EIO] error indication to be lost.
1105	[EISCONN]
1106	Socket is connected. The specified socket is already connected.
1107	[EISDIR]
1108	Is a directory. An attempt was made to open a directory with write mode specified.
1109	[ELOOP]
1110	Symbolic link loop. A loop exists in symbolic links encountered during path name
1111	resolution. This error may also be returned if more than {SYMLOOP_MAX} symbolic links
1112	are encountered during path name resolution.
1113	[EMFILE]
1114	Too many open files. An attempt was made to open more than the maximum number of
1115	{OPEN_MAX} file descriptors allowed in this process.
1116	[EMLINK]
1117	Too many links. An attempt was made to have the link count of a single file exceed
1118	{LINK_MAX}.
1119	[EMSGSIZE]
1120	Message too large. A message sent on a transport provider was larger than an internal
1121	message buffer or some other network limit.
1122	or:
1123	Inappropriate message buffer length.
1124	[EMULTIHOP]
1125	Reserved.
1126	[ENAMETOOLONG]
1127	File name too long. The length of a path name exceeds {PATH_MAX}, or a path name
1128	component is longer than {NAME_MAX}. This error may also occur when path name
1129	substitution, as a result of encountering a symbolic link during path name resolution,
1130	results in a path name string the size of which exceeds {PATH_MAX}.
1131	[ENETDOWN]
1132	Network is down. The local network interface used to reach the destination is down.
1133	[ENETRESET]
1134	The connection was aborted by the network.
1135	[ENETUNREACH]
1136	Network unreachable. No route to the network is present.
1137	[ENFILE]
1138	Too many files open in system. Too many files are currently open in the system. The system
1139	has reached its predefined limit for simultaneously open files and temporarily cannot accept
1140	requests to open another one.

1141	[ENOBUFS]	
1142		No buffer space available. Insufficient buffer resources were available in the system to
1143		perform the socket operation.
1144	XSR [ENODATA]	
1145		No message available. No message is available on the STREAM head read queue.
1146	[ENODEV]	
1147		No such device. An attempt was made to apply an inappropriate function to a device; for
1148		example, trying to read a write-only device such as a printer.
1149	[ENOENT]	
1150		No such file or directory. A component of a specified path name does not exist, or the path
1151		name is an empty string.
1152	[ENOEXEC]	
1153		Executable file format error. A request is made to execute a file that, although it has the
1154		appropriate permissions, is not in the format required by the implementation for executable
1155		files.
1156	[ENOLCK]	
1157		No locks available. A system-imposed limit on the number of simultaneous file and record
1158		locks has been reached and no more are currently available.
1159	[ENOLINK]	
1160		Reserved.
1161	[ENOMEM]	
1162		Not enough space. The new process image requires more memory than is allowed by the
1163		hardware or system-imposed memory management constraints.
1164	[ENOMSG]	
1165		No message of the desired type. The message queue does not contain a message of the
1166		required type during XSI interprocess communication.
1167	[ENOPROTOPT]	
1168		Protocol not available. The protocol option specified to <i>setsockopt()</i> is not supported by the
1169		implementation.
1170	[ENOSPC]	
1171		No space left on a device. During the <i>write()</i> function on a regular file or when extending a
1172		directory, there is no free space left on the device.
1173	XSR [ENOSR]	
1174		No STREAM resources. Insufficient STREAMS memory resources are available to perform a
1175		STREAMS-related function. This is a temporary condition; it may be recovered from if other
1176		processes release resources.
1177	XSR [ENOSTR]	
1178		Not a STREAM. A STREAM function was attempted on a file descriptor that was not
1179		associated with a STREAMS device.
1180	[ENOSYS]	
1181		Function not implemented. An attempt was made to use a function that is not available in
1182		this implementation.
1183	[ENOTCONN]	
1184		Socket not connected. The socket is not connected.

1185	[ENOTDIR]
1186	Not a directory. A component of the specified path name exists, but it is not a directory,
1187	when a directory was expected.
1188	[ENOTEMPTY]
1189	Directory not empty. A directory other than an empty directory was supplied when an
1190	empty directory was expected.
1191	[ENOTSOCK]
1192	Not a socket. The file descriptor does not refer to a socket.
1193	[ENOTSUP]
1194	Not supported. The implementation does not support this feature of the Realtime Option
1195	Group.
1196	[ENOTTY]
1197	Inappropriate I/O control operation. A control function has been attempted for a file or
1198	special file for which the operation is inappropriate.
1199	[ENXIO]
1200	No such device or address. Input or output on a special file refers to a device that does not
1201	exist, or makes a request beyond the capabilities of the device. It may also occur when, for
1202	example, a tape drive is not on-line.
1203	[EOPNOTSUPP]
1204	Operation not supported on socket. The type of socket (address family or protocol) does not
1205	support the requested operation.
1206	[EOVERFLOW]
1207	Value too large to be stored in data type. The user ID or group ID of an IPC or file system
1208	object was too large to be stored into the appropriate member of the caller-provided
1209	structure. This error shall only occur on implementations that support a larger range of user
1210	ID or group ID values than the declared structure member can support. This usually occurs
1211	because the IPC or file system object resides on a remote machine with a larger value of the
1212	type <code>uid_t</code> , <code>off_t</code> , or <code>gid_t</code> than the local system.
1213	[EPERM]
1214	Operation not permitted. An attempt was made to perform an operation limited to
1215	processes with appropriate privileges or to the owner of a file or other resource.
1216	[EPIPE]
1217	Broken pipe. A write was attempted on a socket, pipe, or FIFO for which there is no process
1218	to read the data.
1219	[EPROTO]
1220	Protocol error. Some protocol error occurred. This error is device-specific, but is generally
1221	not related to a hardware failure.
1222	[EPROTONOSUPPORT]
1223	Protocol not supported. The protocol is not supported by the address family, or the protocol
1224	is not supported by the implementation.
1225	[EPROTOTYPE]
1226	Socket type not supported. The socket type is not supported by the protocol.
1227	[ERANGE]
1228	Result too large or too small. The result of the function is too large (overflow) or too small
1229	(underflow) to be represented in the available space (defined in the ISO C standard).

1230	[EROFS]	
1231		Read-only file system. An attempt was made to modify a file or directory on a file system that is read-only.
1232		
1233	[ESPIPE]	
1234		Invalid seek. An attempt was made to access the file offset associated with a pipe or FIFO.
1235	[ESRCH]	
1236		No such process. No process can be found corresponding to that specified by the given process ID.
1237		
1238	[ESTALE]	
1239		Reserved.
1240	XSR [ETIME]	
1241		STREAM <i>ioctl()</i> timeout. The timer set for a STREAMS <i>ioctl()</i> call has expired. The cause of this error is device-specific and could indicate either a hardware or software failure, or a timeout value that is too short for the specific operation. The status of the <i>ioctl()</i> operation is indeterminate.
1242		
1243		
1244		
1245	[ETIMEDOUT]	
1246		Connection timed out. The connection to a remote machine has timed out. If the connection timed out during execution of the function that reported this error (as opposed to timing out prior to the function being called), it is unspecified whether the function has completed some or all of the documented behavior associated with a successful completion of the function.
1247		
1248		
1249		
1250		
1251		or:
1252		Operation timed out. The time limit associated with the operation was exceeded before the operation completed.
1253		
1254	[ETXTBSY]	
1255		Text file busy. An attempt was made to execute a pure-procedure program that is currently open for writing, or an attempt has been made to open for writing a pure-procedure program that is being executed.
1256		
1257		
1258	[EWOULDBLOCK]	
1259		Operation would block. An operation on a socket marked as non-blocking has encountered a situation such as no data available that otherwise would have caused the function to suspend execution.
1260		
1261		
1262		A conforming implementation may assign the same values for [EWOULDBLOCK] and [EAGAIN].
1263		
1264	[EXDEV]	
1265		Improper link. A link to a file on another file system was attempted.
1266	2.3.1 Additional Error Numbers	
1267		Additional implementation-defined error numbers may be defined in <code><errno.h></code> .

1268 **2.4 Signal Concepts**1269 **2.4.1 Signal Generation and Delivery**

1270 A signal is said to be *generated* for (or sent to) a process or thread when the event that causes the
 1271 signal first occurs. Examples of such events include detection of hardware faults, timer
 1272 RTS expiration, signals generated via the **sigevent** structure and terminal activity, as well as
 1273 RTS invocations of *kill()* and *sigqueue()* functions. In some circumstances, the same event generates
 1274 signals for multiple processes.

1275 At the time of generation, a determination is made whether the signal has been generated for the
 1276 process or for a specific thread within the process. Signals which are generated by some action
 1277 attributable to a particular thread, such as a hardware fault, are generated for the thread that
 1278 caused the signal to be generated. Signals that are generated in association with a process ID or
 1279 process group ID or an asynchronous event such as terminal activity are generated for the
 1280 process.

1281 Each process has an action to be taken in response to each signal defined by the system (see
 1282 Section 2.4.3 (on page 530)). A signal is said to be *delivered* to a process when the appropriate
 1283 action for the process and signal is taken. A signal is said to be *accepted* by a process when the
 1284 signal is selected and returned by one of the *sigwait()* functions.

1285 During the time between the generation of a signal and its delivery or acceptance, the signal is
 1286 said to be *pending*. Ordinarily, this interval cannot be detected by an application. However, a
 1287 signal can be *blocked* from delivery to a thread. If the action associated with a blocked signal is
 1288 anything other than to ignore the signal, and if that signal is generated for the thread, the signal
 1289 shall remain pending until it is unblocked, it is accepted when it is selected and returned by a
 1290 call to the *sigwait()* function, or the action associated with it is set to ignore the signal. Signals
 1291 generated for the process shall be delivered to exactly one of those threads within the process
 1292 which is in a call to a *sigwait()* function selecting that signal or has not blocked delivery of the
 1293 signal. If there are no threads in a call to a *sigwait()* function selecting that signal, and if all
 1294 threads within the process block delivery of the signal, the signal shall remain pending on the
 1295 process until a thread calls a *sigwait()* function selecting that signal, a thread unblocks delivery
 1296 of the signal, or the action associated with the signal is set to ignore the signal. If the action
 1297 associated with a blocked signal is to ignore the signal and if that signal is generated for the
 1298 process, it is unspecified whether the signal is discarded immediately upon generation or
 1299 remains pending.

1300 Each thread has a *signal mask* that defines the set of signals currently blocked from delivery to it.
 1301 The signal mask for a thread is initialized from that of its parent or creating thread, or from the
 1302 corresponding thread in the parent process if the thread was created as the result of a call to
 1303 *fork()*. The *sigaction()*, *sigprocmask()*, and *sigsuspend()* functions control the manipulation of the
 1304 signal mask.

1305 The determination of which action is taken in response to a signal is made at the time the signal
 1306 is delivered, allowing for any changes since the time of generation. This determination is
 1307 independent of the means by which the signal was originally generated. If a subsequent
 1308 occurrence of a pending signal is generated, it is implementation-defined as to whether the
 1309 RTS signal is delivered or accepted more than once in circumstances other than those in which
 1310 queuing is required under the Realtime Signals Extension option. The order in which multiple,
 1311 simultaneously pending signals outside the range SIGRTMIN to SIGRTMAX are delivered to or
 1312 accepted by a process is unspecified.

1313 When any stop signal (SIGSTOP, SIGTSTP, SIGTTIN, SIGTTOU) is generated for a process, any
 1314 pending SIGCONT signals for that process shall be discarded. Conversely, when SIGCONT is
 1315 generated for a process, all pending stop signals for that process shall be discarded. When

1316 SIGCONT is generated for a process that is stopped, the process shall be continued, even if the
 1317 SIGCONT signal is blocked or ignored. If SIGCONT is blocked and not ignored, it shall remain
 1318 pending until it is either unblocked or a stop signal is generated for the process.

1319 An implementation shall document any condition not specified by this volume of
 1320 IEEE Std. 1003.1-200x under which the implementation generates signals.

1321 2.4.2 Realtime Signal Generation and Delivery

1322 RTS This section describes extensions to support realtime signal generation and delivery. This
 1323 functionality is dependent on support of the Realtime Signals Extension option (and the rest of
 1324 this section is not further shaded for this option).

1325 Some signal-generating functions, such as high-resolution timer expiration, asynchronous I/O
 1326 completion, interprocess message arrival, and the *sigqueue()* function, support the specification
 1327 of an application-defined value, either explicitly as a parameter to the function or in a **sigevent**
 1328 structure parameter. The **sigevent** structure is defined in <**signal.h**> and shall contain at least
 1329 the following members:

1330

Member Type	Member Name	Description
int	<i>sigev_notify</i>	Notification type.
int	<i>sigev_signo</i>	Signal number.
union signal	<i>sigev_value</i>	Signal value.
void*(unsigned signal)	<i>sigev_notify_function</i>	Notification function.
(pthread_attr_t*)	<i>sigev_notify_attributes</i>	Notification attributes.

1337 The *sigev_notify* member specifies the notification mechanism to use when an asynchronous
 1338 event occurs. This volume of IEEE Std. 1003.1-200x defines the following values for the
 1339 *sigev_notify* member:

1340 SIGEV_NONE No asynchronous notification shall be delivered when the event of
 1341 interest occurs.

1342 SIGEV_SIGNAL The signal specified in *sigev_signo* shall be generated for the process when
 1343 the event of interest occurs. If the implementation supports the Realtime
 1344 Signals Extension option and if the SA_SIGINFO flag is set for that signal
 1345 number, then the signal shall be queued to the process and the value
 1346 specified in *sigev_value* shall be the *si_value* component of the generated
 1347 signal. If SA_SIGINFO is not set for that signal number, it is unspecified
 1348 whether the signal is queued and what value, if any, is sent.

1349 SIGEV_THREAD A notification function shall be called to perform notification.

1350 An implementation may define additional notification mechanisms.

1351 The *sigev_signo* member specifies the signal to be generated. The *sigev_value* member is the
 1352 application-defined value to be passed to the signal-catching function at the time of the signal
 1353 delivery or to be returned at signal acceptance as the *si_value* member of the **siginfo_t** structure.

1354 The **signal** union is defined in <**signal.h**> and contains at least the following members:

1355
1356
1357
1358

Member Type	Member Name	Description
int	<i>sival_int</i>	Integer signal value.
void*	<i>sival_ptr</i>	Pointer signal value.

1359
1360

The *sival_int* member is used when the application-defined value is of type **int**; the *sival_ptr* member is used when the application-defined value is a pointer.

1361
1362
1363
1364
1365
1366

When a signal is generated by the *sigqueue()* function or any signal-generating function that supports the specification of an application-defined value, the signal shall be marked pending and, if the SA_SIGINFO flag is set for that signal, the signal shall be queued to the process along with the application-specified signal value. Multiple occurrences of signals so generated are queued in FIFO order. It is unspecified whether signals so generated are queued when the SA_SIGINFO flag is not set for that signal.

1367
1368
1369
1370

Signals generated by the *kill()* function or other events that cause signals to occur, such as detection of hardware faults, *alarm()* timer expiration, or terminal activity, and for which the implementation does not support queuing, have no effect on signals already queued for the same signal number.

1371
1372
1373

When multiple unblocked signals, all in the range SIGRTMIN to SIGRTMAX, are pending, the behavior shall be as if the implementation delivers the pending unblocked signal with the lowest signal number within that range. No other ordering of signal delivery is specified.

1374
1375

If, when a pending signal is delivered, there are additional signals queued to that signal number, the signal remains pending. Otherwise, the pending indication is reset.

1376
1377
1378

Multi-threaded programs can use an alternate event notification mechanism. When a notification is processed, and the *sigev_notify* member of the **sigevent** structure has the value SIGEV_THREAD, the function *sigev_notify_function* is called with parameter *sigev_value*.

1379
1380
1381
1382
1383
1384

The function shall be executed in an environment as if it were the *start_routine* for a newly created thread with thread attributes specified by *sigev_notify_attributes*. If *sigev_notify_attributes* is NULL, the behavior shall be as if the thread were created with the *detachstate* attribute set to PTHREAD_CREATE_DETACHED. Supplying an attributes structure with a *detachstate* attribute of PTHREAD_CREATE_JOINABLE results in undefined behavior. The signal mask of this thread is implementation-defined.

1385 2.4.3 Signal Actions

1386
1387
1388

There are three types of action that can be associated with a signal: SIG_DFL, SIG_IGN, or a pointer to a function. Initially, all signals shall be set to SIG_DFL or SIG_IGN prior to entry of the *main()* routine (see the *exec* functions). The actions prescribed by these values are as follows:

1389

SIG_DFL Signal-specific default action.

1390
1391 RTS
1392
1393

The default actions for the signals defined in this volume of IEEE Std. 1003.1-200x are specified under **<signal.h>**. If the Realtime Signals Extension option is supported, the default actions for the realtime signals in the range SIGRTMIN to SIGRTMAX are to terminate the process abnormally.

1394
1395
1396
1397
1398
1399
1400

If the default action is to stop the process, the execution of that process is temporarily suspended. When a process stops, a SIGCHLD signal shall be generated for its parent process, unless the parent process has set the SA_NOCLDSTOP flag. While a process is stopped, any additional signals that are sent to the process shall not be delivered until the process is continued, except SIGKILL which always terminates the receiving process. A process that is a member of an orphaned process group shall not be allowed to stop in response to

1401 the SIGTSTP, SIGTTIN, or SIGTTOU signals. In cases where delivery of one of
 1402 these signals would stop such a process, the signal shall be discarded.

1403 Setting a signal action to SIG_DFL for a signal that is pending, and whose default
 1404 action is to ignore the signal (for example, SIGCHLD), shall cause the pending
 1405 signal to be discarded, whether or not it is blocked.

1406 The default action for SIGCONT is to resume execution at the point where the
 1407 RTS process was stopped, after first handling any pending unblocked signals. If the
 1408 Realttime Signals Extension option is supported, any queued values pending shall
 1409 be discarded and the resources used to queue them shall be released and returned
 1410 to the system for other use.

1411 SIG_IGN Ignore signal.

1412 Delivery of the signal shall have no effect on the process. The behavior of a process
 1413 RTS is undefined after it ignores a SIGFPE, SIGILL, SIGSEGV, or SIGBUS signal that
 1414 RTS was not generated by *kill()*, *sigqueue()*, or *raise()*.

1415 The system shall not allow the action for the signals SIGKILL or SIGSTOP to be set
 1416 to SIG_IGN.

1417 Setting a signal action to SIG_IGN for a signal that is pending shall cause the
 1418 pending signal to be discarded, whether or not it is blocked.

1419 If a process sets the action for the SIGCHLD signal to SIG_IGN, the behavior is
 1420 XSI unspecified, except as specified below.

1421 If the action for the SIGCHLD signal is set to SIG_IGN, child processes of the
 1422 calling processes shall not be transformed into zombie processes when they
 1423 terminate. If the calling process subsequently waits for its children, and the process
 1424 has no unwaited-for children that were transformed into zombie processes, it shall
 1425 block until all of its children terminate, and *wait()*, *waitid()*, and *waitpid()* shall fail
 1426 and set *errno* to [ECHILD].

1427 RTS If the Realttime Signals Extension option is supported, any queued values pending
 1428 shall be discarded and the resources used to queue them shall be released and
 1429 made available to queue other signals.

1430 *pointer to a function*

1431 Catch signal.

1432 On delivery of the signal, the receiving process is to execute the signal-catching
 1433 function at the specified address. After returning from the signal-catching function,
 1434 the receiving process shall resume execution at the point at which it was
 1435 interrupted.

1436 If the SA_SIGINFO flag for the signal is cleared, the signal-catching function shall
 1437 be entered as a C-language function call as follows:

1438 `void func(int signo);`

1439 XSI|RTS If the SA_SIGINFO flag for the signal is set, the signal-catching function shall be
 1440 entered as a C-language function call as follows:

1441 `void func(int signo, siginfo_t *info, void *context);`

1442 where *func* is the specified signal-catching function, *signo* is the signal number of
 1443 the signal being delivered, and *info* is a pointer to a **siginfo_t** structure defined in
 1444 **<signal.h>** containing at least the following members:

1445
1446
1447
1448
1449

	Member Type	Member Name	Description
	int	<i>si_signo</i>	Signal number
	int	<i>si_code</i>	Cause of the signal
	union signal	<i>si_value</i>	Signal value

1450
1451
1452

The *si_signo* member contains the signal number. This is the same as the *signo* parameter. The *si_code* member contains a code identifying the cause of the signal. The following values are defined for *si_code*:

1453 **Notes to Reviewers**

1454
1455

This section with side shading will not appear in the final copy. - Ed.

The shading in this area needs some work.

1456 XSI|RTS
1457
1458
1459

SI_USER The signal was sent by the *kill()* function. The implementation may set *si_code* to **SI_USER** if the signal was sent by the *raise()* or *abort()* functions or any similar functions provided as implementation extensions.

1460 RTS

SI_QUEUE The signal was sent by the *sigqueue()* function.

1461 RTS
1462

SI_TIMER The signal was generated by the expiration of a timer set by *timer_settime()*.

1463 RTS
1464

SI_ASYNCIO The signal was generated by the completion of an asynchronous I/O request.

1465 RTS
1466

SI_MESGQ The signal was generated by the arrival of a message on an empty message queue.

1467
1468
1469

If the signal was not generated by one of the functions or events listed above, the *si_code* shall be set to an implementation-defined value that is not equal to any of the values defined above.

1470 RTS
1471
1472

If the Realtime Signals Extension is supported, and *si_code* is one of **SI_QUEUE**, **SI_TIMER**, **SI_ASYNCIO**, or **SI_MESGQ**, then *si_value* contains the application-specified signal value. Otherwise, the contents of *si_value* are undefined.

1473
1474 XSI
1475 RTS

The behavior of a process is undefined after it returns normally from a signal-catching function for a **SIGBUS**, **SIGFPE**, **SIGILL**, or **SIGSEGV** signal that was not generated by *kill()*, *sigqueue()*, or *raise()*.

1476

The system shall not allow a process to catch the signals **SIGKILL** and **SIGSTOP**.

1477
1478
1479

If a process establishes a signal-catching function for the **SIGCHLD** signal while it has a terminated child process for which it has not waited, it is unspecified whether a **SIGCHLD** signal is generated to indicate that child process.

1480
1481
1482
1483

When signal-catching functions are invoked asynchronously with process execution, the behavior of some of the functions defined by this volume of IEEE Std. 1003.1-200x is unspecified if they are called from a signal-catching function.

1484
1485
1486

The following table defines a set of functions that are either reentrant or not interruptible by signals and are async-signal-safe. Therefore applications may invoke them, without restriction, from signal-catching functions:

1487 **Notes to Reviewers**1488 *This section with side shading will not appear in the final copy. - Ed.*

1489 The contents of the following tables need to be reviewed.

1490 Base functions:

1491	<code>_Exit()</code>	<code>fpathconf()</code>	<code>pipe()</code>	<code>stat()</code>
1492	<code>_exit()</code>	<code>fstat()</code>	<code>raise()</code>	<code>symlink()</code>
1493	<code>access()</code>	<code>fsync()</code>	<code>read()</code>	<code>sysconf()</code>
1494	<code>alarm()</code>	<code>ftruncate()</code>	<code>readlink()</code>	<code>tcdrain()</code>
1495	<code>cfgetispeed()</code>	<code>getegid()</code>	<code>rename()</code>	<code>tcflow()</code>
1496	<code>cfgetospeed()</code>	<code>geteuid()</code>	<code>rmdir()</code>	<code>tcflush()</code>
1497	<code>cfsetispeed()</code>	<code>getgid()</code>	<code>setgid()</code>	<code>tcgetattr()</code>
1498	<code>cfsetospeed()</code>	<code>getgroups()</code>	<code>setpgid()</code>	<code>tcsetpgrp()</code>
1499	<code>chdir()</code>	<code>getpgrp()</code>	<code>setsid()</code>	<code>tcsendbreak()</code>
1500	<code>chmod()</code>	<code>getpid()</code>	<code>setuid()</code>	<code>tcsetattr()</code>
1501	<code>chown()</code>	<code>getppid()</code>	<code>sigaction()</code>	<code>tcsetpgrp()</code>
1502	<code>close()</code>	<code>getuid()</code>	<code>sigaddset()</code>	<code>time()</code>
1503	<code>creat()</code>	<code>kill()</code>	<code>sigdelset()</code>	<code>times()</code>
1504	<code>dup()</code>	<code>link()</code>	<code>sigemptyset()</code>	<code>umask()</code>
1505	<code>dup2()</code>	<code>lseek()</code>	<code>sigfillset()</code>	<code>uname()</code>
1506	<code>execle()</code>	<code>lstat()</code>	<code>sigismember()</code>	<code>unlink()</code>
1507	<code>execve()</code>	<code>mkdir()</code>	<code>signal()</code>	<code>utime()</code>
1508	<code>fchmod()</code>	<code>mkfifo()</code>	<code>sigpending()</code>	<code>wait()</code>
1509	<code>fchown()</code>	<code>open()</code>	<code>sigprocmask()</code>	<code>waitpid()</code>
1510	<code>fcntl()</code>	<code>pathconf()</code>	<code>sigsuspend()</code>	<code>write()</code>
1511	<code>fork()</code>	<code>pause()</code>	<code>sleep()</code>	

1512 Realtime functions:

1513	<code>aio_error()</code>	<code>clock_gettime()</code>	<code>sigpause()</code>	<code>timer_getoverrun()</code>
1514	<code>aio_return()</code>	<code>fdatasync()</code>	<code>sigqueue()</code>	<code>timer_gettime()</code>
1515	<code>aio_suspend()</code>	<code>sem_post()</code>	<code>sigset()</code>	<code>timer_settime()</code>

1516 Tracing functions:

1517 `posix_trace_event()` |

1518 All functions not in the above table are considered to be unsafe with respect to |
 1519 signals. In the presence of signals, all functions defined by this volume of |
 1520 IEEE Std. 1003.1-200x shall behave as defined when called from or interrupted by a |
 1521 signal-catching function, with a single exception: when a signal interrupts an |
 1522 unsafe function and the signal-catching function calls an unsafe function, the |
 1523 behavior is undefined.

1524 When a signal is delivered to a thread, if the action of that signal specifies termination, stop, or
 1525 continue, the entire process shall be terminated, stopped, or continued, respectively.

1526 **2.4.4 Signal Effects on Other Functions**

1527 Signals affect the behavior of certain functions defined by this volume of IEEE Std. 1003.1-200x if
1528 delivered to a process while it is executing such a function. If the action of the signal is to
1529 terminate the process, the process shall be terminated and the function shall not return. If the
1530 action of the signal is to stop the process, the process shall stop until continued or terminated.
1531 Generation of a SIGCONT signal for the process shall cause the process to be continued, and the
1532 original function shall continue at the point the process was stopped. If the action of the signal is
1533 to invoke a signal-catching function, the signal-catching function shall be invoked; in this case
1534 the original function is said to be *interrupted* by the signal.

1535 **Notes to Reviewers**

1536 *This section with side shading will not appear in the final copy. - Ed.*

1537 D3, XSH, ERN 20 points out a discrepancy between the following sentence and the paragraph
1538 above beginning "All functions not in the above ...". An interpretation will be filed.
1539 If the signal-catching function executes a **return** statement, the behavior of the interrupted
1540 function shall be as described individually for that function. Signals that are ignored shall not
1541 affect the behavior of any function; signals that are blocked shall not affect the behavior of any
1542 function until they are unblocked and then delivered, except as specified for the *sigpending()* and
1543 *sigwait()* functions.

1544 2.5 Standard I/O Streams

1545 A stream is associated with an external file (which may be a physical device) by *opening* a file,
1546 which may involve *creating* a new file. Creating an existing file causes its former contents to be
1547 discarded if necessary. If a file can support positioning requests, (such as a disk file, as opposed
1548 to a terminal), then a *file position indicator* associated with the stream is positioned at the start
1549 (byte number 0) of the file, unless the file is opened with append mode, in which case it is
1550 implementation-defined whether the file position indicator is initially positioned at the
1551 beginning or end of the file. The file position indicator is maintained by subsequent reads,
1552 writes, and positioning requests, to facilitate an orderly progression through the file. All input
1553 takes place as if bytes were read by successive calls to *fgetc()*; all output takes place as if bytes
1554 were written by successive calls to *fputc()*.

1555 When a stream is *unbuffered*, bytes are intended to appear from the source or at the destination
1556 as soon as possible; otherwise, bytes may be accumulated and transmitted as a block. When a
1557 stream is *fully buffered*, bytes are intended to be transmitted as a block when a buffer is filled.
1558 When a stream is *line buffered*, bytes are intended to be transmitted as a block when a newline
1559 byte is encountered. Furthermore, bytes are intended to be transmitted as a block when a buffer
1560 is filled, when input is requested on an unbuffered stream, or when input is requested on a line-
1561 buffered stream that requires the transmission of bytes. Support for these characteristics is
1562 implementation-defined, and may be affected via *setbuf()* and *setvbuf()*.

1563 A file may be disassociated from a controlling stream by *closing* the file. Output streams are
1564 flushed (any unwritten buffer contents are transmitted) before the stream is disassociated from
1565 the file. The value of a pointer to a FILE object is indeterminate after the associated file is closed
1566 (including the standard streams).

1567 A file may be subsequently reopened, by the same or another program execution, and its
1568 contents reclaimed or modified (if it can be repositioned at its start). If the *main()* function
1569 returns to its original caller, or if the *exit()* function is called, all open files are closed (hence all
1570 output streams are flushed) before program termination. Other paths to program termination,
1571 such as calling *abort()*, need not close all files properly.

1572 The address of the FILE object used to control a stream may be significant; a copy of a FILE
1573 object need not necessarily serve in place of the original.

1574 At program start-up, three streams are predefined and need not be opened explicitly: *standard*
1575 *input* (for reading conventional input), *standard output* (for writing conventional output), and
1576 *standard error* (for writing diagnostic output). When opened, the standard error stream is not
1577 fully buffered; the standard input and standard output streams are fully buffered if and only if
1578 the stream can be determined not to refer to an interactive device.

1579 2.5.1 Interaction of File Descriptors and Standard I/O Streams

1580 cx This section describes the interaction of file descriptors and standard I/O streams. This
1581 functionality is an extension to the ISO C standard (and the rest of this section is not further CX
1582 shaded).

1583 An open file description may be accessed through a file descriptor, which is created using
1584 functions such as *open()* or *pipe()*, or through a stream, which is created using functions such as
1585 *fopen()* or *popen()*. Either a file descriptor or a stream is called a *handle* on the open file
1586 description to which it refers; an open file description may have several handles.

1587 Handles can be created or destroyed by explicit user action, without affecting the underlying
1588 open file description. Some of the ways to create them include *fcntl()*, *dup()*, *fdopen()*, *fileno()*,
1589 and *fork()*. They can be destroyed by at least *fclose()*, *close()*, and the *exec* functions.

1590 A file descriptor that is never used in an operation that could affect the file offset (for example,
 1591 *read()*, *write()*, or *lseek()*) is not considered a handle for this discussion, but could give rise to one
 1592 (for example, as a consequence of *fdopen()*, *dup()*, or *fork()*). This exception does not include the
 1593 file descriptor underlying a stream, whether created with *fopen()* or *fdopen()*, so long as it is not
 1594 used directly by the application to affect the file offset. The *read()* and *write()* functions
 1595 implicitly affect the file offset; *lseek()* explicitly affects it.

1596 The result of function calls involving any one handle (the *active handle*) is defined elsewhere in
 1597 this volume of IEEE Std. 1003.1-200x, but if two or more handles are used, and any one of them is
 1598 a stream, the application shall ensure that their actions are coordinated as described below. If
 1599 this is not done, the result is undefined.

1600 A handle which is a stream is considered to be closed when either an *fclose()* or *freopen()* is
 1601 executed on it (the result of *freopen()* is a new stream, which cannot be a handle on the same
 1602 open file description as its previous value), or when the process owning that stream terminates
 1603 with *exit()* or *abort()*. A file descriptor is closed by *close()*, *_exit()*, or the *exec* functions when
 1604 FD_CLOEXEC is set on that file descriptor.

1605 For a handle to become the active handle, the application shall ensure that the actions below are
 1606 performed between the last use of the handle (the current active handle) and the first use of the
 1607 second handle (the future active handle). The second handle then becomes the active handle. All
 1608 activity by the application affecting the file offset on the first handle shall be suspended until it
 1609 again becomes the active file handle. (If a stream function has as an underlying function one that
 1610 affects the file offset, the stream function shall be considered to affect the file offset.)

1611 The handles need not be in the same process for these rules to apply.

1612 Note that after a *fork()*, two handles exist where one existed before. The application shall ensure
 1613 that, if both handles can ever be accessed, they are both in a state where the other could become
 1614 the active handle first. The application shall prepare for a *fork()* exactly as if it were a change of
 1615 active handle. (If the only action performed by one of the processes is one of the *exec* functions or
 1616 *_exit()* (not *exit()*), the handle is never accessed in that process.)

1617 For the first handle, the first applicable condition below applies. After the actions required
 1618 below are taken, if the handle is still open, the application can close it.

- 1619 • If it is a file descriptor, no action is required.
- 1620 • If the only further action to be performed on any handle to this open file descriptor is to close
 1621 it, no action need be taken.
- 1622 • If it is a stream which is unbuffered, no action need be taken.
- 1623 • If it is a stream which is line buffered, and the last byte written to the stream was a newline
 1624 (that is, as if a:
 1625

```
putc( '\n' )
```


 1626 was the most recent operation on that stream), no action need be taken.
- 1627 • If it is a stream which is open for writing or appending (but not also open for reading), the
 1628 application shall either perform an *flush()*, or the stream shall be closed.
- 1629 • If the stream is open for reading and it is at the end of the file (*feof()* is true), no action need
 1630 be taken.
- 1631 • If the stream is open with a mode that allows reading and the underlying open file
 1632 description refers to a device that is capable of seeking, the application shall either perform
 1633 an *flush()*, or the stream shall be closed.

1634 Otherwise, the result is undefined.

1635 For the second handle:

- 1636 • If any previous active handle has been used by a function that explicitly changed the file
- 1637 offset, except as required above for the first handle, the application shall perform an *lseek()* or
- 1638 *fseek()* (as appropriate to the type of handle) to an appropriate location.

1639 If the active handle ceases to be accessible before the requirements on the first handle, above,

1640 have been met, the state of the open file description becomes undefined. This might occur during

1641 functions such as a *fork()* or *_exit()*.

1642 The *exec* functions make inaccessible all streams that are open at the time they are called,

1643 independent of which streams or file descriptors may be available to the new process image.

1644 When these rules are followed, regardless of the sequence of handles used, implementations

1645 shall ensure that an application, even one consisting of several processes, shall yield correct

1646 results: no data shall be lost or duplicated when writing, and all data shall be written in order,

1647 except as requested by seeks. It is implementation-defined whether, and under what conditions,

1648 all input is seen exactly once.

1649 If the rules above are not followed, the result is unspecified.

1650 Each function that operates on a stream is said to have zero or more *underlying functions*. This

1651 means that the stream function shares certain traits with the underlying functions, but does not

1652 require that there be any relation between the implementations of the stream function and its

1653 underlying functions.

1654 2.5.2 Stream Orientation and Encoding Rules

1655 For conformance to the ISO/IEC 9899:1999 standard, the definition of a stream includes an

1656 *orientation*. After a stream is associated with an external file, but before any operations are

1657 performed on it, the stream is without orientation. Once a wide-character input/output function

1658 has been applied to a stream without orientation, the stream shall become *wide-oriented*.

1659 Similarly, once a byte input/output function has been applied to a stream without orientation,

1660 the stream shall become *byte-oriented*. Only a call to the *freopen()* function or the *fwide()* function

1661 can otherwise alter the orientation of a stream.

1662 A successful call to *freopen()* shall remove any orientation. The three predefined streams *standard*

1663 *input*, *standard output*, and *standard error* shall be unoriented at program start-up.

1664 Byte input/output functions cannot be applied to a wide-oriented stream, and wide-character

1665 input/output functions cannot be applied to a byte-oriented stream. The remaining stream

1666 operations shall not affect and shall not be affected by a stream's orientation, except for the

1667 following additional restrictions:

- 1668 • Binary wide-oriented streams have the file positioning restrictions ascribed to both text and
- 1669 binary streams.
- 1670 • For wide-oriented streams, after a successful call to a file-positioning function that leaves the
- 1671 file position indicator prior to the end-of-file, a wide-character output function can overwrite
- 1672 a partial character; any file contents beyond the byte(s) written are henceforth undefined.

1673 Each wide-oriented stream has an associated **mbstate_t** object that stores the current parse state

1674 of the stream. A successful call to *fgetpos()* shall store a representation of the value of this

1675 **mbstate_t** object as part of the value of the **fpos_t** object. A later successful call to *fsetpos()* using

1676 the same stored **fpos_t** value shall restore the value of the associated **mbstate_t** object as well as

1677 the position within the controlled stream.

1678 Implementations that support multiple encoding rules associate an encoding rule with the
1679 stream. The encoding rule shall be determined by the setting of the *LC_CTYPE* category in the
1680 current locale at the time when the stream becomes wide-oriented. If a wide-character
1681 input/output function is applied to a byte-oriented stream, the encoding rule used is undefined.
1682 As with the stream's orientation, the encoding rule associated with a stream cannot be changed
1683 once it has been set, except by a successful call to *freopen()* which clears the encoding rule and
1684 resets the orientation to unoriented.

1685 Although both text and binary wide-oriented streams are conceptually sequences of wide
1686 characters, the external file associated with a wide-oriented stream is a sequence of (possibly
1687 multibyte) characters generalized as follows:

- 1688 • Multibyte encodings within files may contain embedded null bytes (unlike multibyte
1689 encodings valid for use internal to the program).
- 1690 • A file need not begin nor end in the initial shift state.

1691 Moreover, the encodings used for characters may differ among files. Both the nature and choice
1692 of such encodings are implementation-defined.

1693 The wide-character input functions read characters from the stream and convert them to wide
1694 characters as if they were read by successive calls to the *fgetwc()* function. Each conversion shall
1695 occur as if by a call to the *mbrtowc()* function, with the conversion state described by the stream's
1696 own **mbstate_t** object, except the encoding rule associated with the stream is used instead of the
1697 encoding rule implied by the *LC_CTYPE* category of the current locale.

1698 The wide-character output functions convert wide characters to (possibly multibyte) characters
1699 and write them to the stream as if they were written by successive calls to the *fputwc()* function.
1700 Each conversion shall occur as if by a call to the *wcrtomb()* function, with the conversion state
1701 described by the stream's own **mbstate_t** object, except the encoding rule associated with the
1702 stream is used instead of the encoding rule implied by the *LC_CTYPE* category of the current
1703 locale.

1704 An *encoding error* shall occur if the character sequence presented to the underlying *mbrtowc()*
1705 function does not form a valid (generalized) character, or if the code value passed to the
1706 underlying *wcrtomb()* function does not correspond to a valid (generalized) character. The
1707 wide-character input/output functions and the byte input/output functions store the value of
1708 the macro *EILSEQ* in *errno* if and only if an encoding error occurs.

1709 **2.6 STREAMS**

1710 XSR STREAMS functionality is provided on implementations supporting the XSI STREAMS Option
 1711 Group. This functionality is dependent on support of the XSI STREAMS option (and the rest of
 1712 this section is not further shaded for this option).

1713 STREAMS provides a uniform mechanism for implementing networking services and other
 1714 character-based I/O. The STREAMS function provides direct access to protocol modules. A
 1715 STREAM is typically a full-duplex connection between a process and an open device or pseudo-
 1716 device. However, since pipes may be STREAMS-based, a STREAM can be a full-duplex
 1717 connection between two processes. The STREAM itself exists entirely within the implementation
 1718 and provides a general character I/O function for processes. It optionally includes one or more
 1719 intermediate processing modules that are interposed between the process end of the STREAM
 1720 (STREAM head) and a device driver at the end of the STREAM (STREAM end).

1721 STREAMS I/O is based on messages. Messages flow in both directions in a STREAM. A given
 1722 module need not understand and process every message in the STREAM, but every module in
 1723 the STREAM handles every message. Each module accepts messages from one of its neighbor
 1724 modules in the STREAM, and passes them to the other neighbor. For example, a line discipline
 1725 module may transform the data. Data flow through the intermediate modules is bidirectional,
 1726 with all modules handling, and optionally processing, all messages. There are three types of
 1727 messages:

- 1728 • *Data messages* containing actual data for input or output
- 1729 • *Control data* containing instructions for the STREAMS modules and underlying
 1730 implementation
- 1731 • Other messages, which include file descriptors

1732 The function between the STREAM and the rest of the implementation is provided by a set of
 1733 functions at the STREAM head. When a process calls *write()*, *putmsg()*, *putpmsg()*, or *ioctl()*,
 1734 messages are sent down the STREAM, and *read()*, *getmsg()*, or *getpmsg()* accepts data from the
 1735 STREAM and passes it to a process. Data intended for the device at the downstream end of the
 1736 STREAM is packaged into messages and sent downstream, while data and signals from the
 1737 device are composed into messages by the device driver and sent upstream to the STREAM
 1738 head.

1739 When a STREAMS-based device is opened, a STREAM is created that contains two modules: the
 1740 STREAM head module and the STREAM end (driver) module. If pipes are STREAMS-based in
 1741 an implementation, when a pipe is created, two STREAMS are created, each containing a
 1742 STREAM head module. Other modules are added to the STREAM using *ioctl()*. New modules
 1743 are “pushed” onto the STREAM one at a time in last-in, first-out (LIFO) style, as though the
 1744 STREAM was a push-down stack.

1745 **Priority**

1746 Message types are classified according to their queuing priority and may be *normal* (non-
 1747 priority), *priority*, or *high-priority* messages. A message belongs to a particular priority band that
 1748 determines its ordering when placed on a queue. Normal messages have a priority band of 0 and
 1749 are always placed at the end of the queue following all other messages in the queue. High-
 1750 priority messages are always placed at the head of a queue, but are discarded if there is already a
 1751 high-priority message in the queue. Their priority band is ignored; they are high-priority by
 1752 virtue of their type. Priority messages have a priority band greater than 0. Priority messages are
 1753 always placed after any messages of the same or higher priority. High-priority and priority
 1754 messages are used to send control and data information outside the normal flow of control. By
 1755 convention, high-priority messages are not affected by flow control. Normal and priority

1756 messages have separate flow controls.

1757 **Message Parts**

1758 A process may access STREAMS messages that contain a data part, control part, or both. The
1759 data part is that information which is transmitted over the communication medium and the
1760 control information is used by the local STREAMS modules. The other types of messages are
1761 used between modules and are not accessible to processes. Messages containing only a data part
1762 are accessible via *putmsg()*, *putpmsg()*, *getmsg()*, *getpmsg()*, *read()*, or *write()*. Messages
1763 containing a control part with or without a data part are accessible via calls to *putmsg()*,
1764 *putpmsg()*, *getmsg()*, or *getpmsg()*.

1765 **2.6.1 Accessing STREAMS**

1766 A process accesses STREAMS-based files using the standard functions *close()*, *ioctl()*, *getmsg()*,
1767 *getpmsg()*, *open()*, *pipe()*, *poll()*, *putmsg()*, *putpmsg()*, *read()*, or *write()*. Refer to the applicable
1768 function definitions for general properties and errors.

1769 Calls to *ioctl()* are used to perform control functions with the STREAMS-based device associated
1770 with the file descriptor *fildes*. The arguments *command* and *arg* are passed to the STREAMS file
1771 designated by *fildes* and are interpreted by the STREAM head. Certain combinations of these
1772 arguments may be passed to a module or driver in the STREAM.

1773 Since these STREAMS requests are a subset of *ioctl()*, they are subject to the errors described
1774 there.

1775 STREAMS modules and drivers can detect errors, sending an error message to the STREAM
1776 head, thus causing subsequent functions to fail and set *errno* to the value specified in the
1777 message. In addition, STREAMS modules and drivers can elect to fail a particular *ioctl()* request
1778 alone by sending a negative acknowledgement message to the STREAM head. This causes just
1779 the pending *ioctl()* request to fail and set *errno* to the value specified in the message.

1780 2.7 XSI Interprocess Communication

1781 XSI This section describes extensions to support interprocess communication. This functionality is
 1782 dependent on support of the XSI Extension (and the rest of this section is not further shaded for
 1783 this option).

1784 The following message passing, semaphore, and shared memory services form an XSI
 1785 interprocess communication facility. Certain aspects of their operation are common, and are
 1786 described below.

1787
 1788

IPC Functions		
<i>msgctl()</i>	<i>semctl()</i>	<i>shmctl()</i>
<i>msgget()</i>	<i>semget()</i>	<i>shmdt()</i>
<i>msgrcv()</i>	<i>semop()</i>	<i>shmget()</i>
<i>msgsnd()</i>	<i>shmat()</i>	

1789
 1790
 1791
 1792

1793 Another interprocess communication facility is provided by functions in the Realtime Option
 1794 Group; see Section 2.8 (on page 543).

1795 2.7.1 IPC General Description

1796 Each individual shared memory segment, message queue, and semaphore set is identified by a
 1797 unique positive integer, called respectively a shared memory identifier, *shmid*, a semaphore
 1798 identifier, *semid*, and a message queue identifier, *msqid*. The identifiers are returned by calls to
 1799 *shmget()*, *semget()*, and *msgget()*, respectively.

1800 Associated with each identifier is a data structure which contains data related to the operations
 1801 which may be or may have been performed; see the Base Definitions volume of
 1802 IEEE Std. 1003.1-200x, <sys/shm.h>, <sys/sem.h>, and <sys/msg.h> for their descriptions.

1803 Each of the data structures contains both ownership information and an **ipc_perm** structure (see
 1804 the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/ipc.h>) which are used in conjunction
 1805 to determine whether or not read/write (read/alter for semaphores) permissions should be
 1806 granted to processes using the IPC facilities. The *mode* member of the **ipc_perm** structure acts as
 1807 a bit field which determines the permissions.

1808 The values of the bits are given below in octal notation.

1809
 1810

Bit	Meaning
0400	Read by user.
0200	Write by user.
0040	Read by group.
0020	Write by group.
0004	Read by others.
0002	Write by others.

1811
 1812
 1813
 1814
 1815
 1816

1817 The name of the **ipc_perm** structure is *shm_perm*, *sem_perm*, or *msg_perm*, depending on which
 1818 service is being used. In each case, read and write/alter permissions are granted to a process if
 1819 one or more of the following are true ("xxx" is replaced by *shm*, *sem*, or *msg*, as appropriate):

- 1820 • The process has appropriate privileges.
- 1821 • The effective user ID of the process matches *xxx_perm.cuid* or *xxx_perm.uid* in the data
 1822 structure associated with the IPC identifier, and the appropriate bit of the *user* field in
 1823 *xxx_perm.mode* is set.

- 1824
- 1825
- 1826
- 1827
- The effective user ID of the process does not match *xxx_perm.cuid* or *xxx_perm.uid* but the effective group ID of the process matches *xxx_perm.cgid* or *xxx_perm.gid* in the data structure associated with the IPC identifier, and the appropriate bit of the *group* field in *xxx_perm.mode* is set.
- 1828
- 1829
- 1830
- 1831
- The effective user ID of the process does not match *xxx_perm.cuid* or *xxx_perm.uid* and the effective group ID of the process does not match *xxx_perm.cgid* or *xxx_perm.gid* in the data structure associated with the IPC identifier, but the appropriate bit of the *other* field in *xxx_perm.mode* is set.
- 1832
- Otherwise, the permission is denied.

1833 2.8 Realtime

1834 This section defines functions to support the source portability of applications with realtime
1835 requirements. The presence of many of these functions is dependent on support for
1836 implementation options described in the text.

1837 The specific functional areas included in this section and their scope include the following. Full
1838 definitions of these terms can be found in the Base Definitions volume of IEEE Std. 1003.1-200x,
1839 Chapter 3, Definitions.

- 1840 • Semaphores
- 1841 • Process Memory Locking
- 1842 • Memory Mapped Files and Shared Memory Objects
- 1843 • Priority Scheduling
- 1844 • Realtime Signal Extension
- 1845 • Timers
- 1846 • Interprocess Communication
- 1847 • Synchronized Input and Output
- 1848 • Asynchronous Input and Output

1849 All the realtime functions defined in this volume of IEEE Std. 1003.1-200x are portable, although
1850 some of the numeric parameters used by an implementation may have hardware dependencies.

1851 2.8.1 Realtime Signals

1852 RTS Realtime signal generation and delivery is dependent on support for the Realtime Signals
1853 Extension option.

1854 See Section 2.4.2 (on page 529).

1855 2.8.2 Asynchronous I/O

1856 AIO The functionality described in this section is dependent on support of the Asynchronous Input
1857 and Output option (and the rest of this section is not further shaded for this option).

1858 An asynchronous I/O control block structure **aiocb** is used in many asynchronous I/O
1859 functions. It is defined in the Base Definitions volume of IEEE Std. 1003.1-200x, <**aio.h**> and has
1860 at least the following members:

1861	Member Type	Member Name	Description
1862	int	<i>aio_fildes</i>	File descriptor.
1863	off_t	<i>aio_offset</i>	File offset.
1864	volatile void*	<i>aio_buf</i>	Location of buffer.
1865	size_t	<i>aio_nbytes</i>	Length of transfer.
1866	int	<i>aio_reqprio</i>	Request priority offset.
1867	struct sigevent	<i>aio_sigevent</i>	Signal number and value.
1868	int	<i>aio_lio_opcode</i>	Operation to be performed.

1869 The *aio_fildes* element is the file descriptor on which the asynchronous operation is performed.

1870 If **O_APPEND** is not set for the file descriptor *aio_fildes* and if *aio_fildes* is associated with a
1871 device that is capable of seeking, then the requested operation takes place at the absolute
1872 position in the file as given by *aio_offset*, as if *lseek()* were called immediately prior to the

1873 operation with an *offset* argument equal to *aio_offset* and a *whence* argument equal to `SEEK_SET`.
 1874 If `O_APPEND` is set for the file descriptor, or if *aio_fildes* is associated with a device that is
 1875 incapable of seeking, write operations append to the file in the same order as the calls were
 1876 made, with the following exception: under implementation-defined circumstances, such as
 1877 operation on a multiprocessor or when requests of differing priorities are submitted at the same
 1878 time, the ordering restriction may be relaxed. After a successful call to enqueue an asynchronous
 1879 I/O operation, the value of the file offset for the file is unspecified. The *aio_nbytes* and *aio_buf*
 1880 elements are the same as the *nbyte* and *buf* arguments defined by `read()` and `write()`, respectively.

1881 If `_POSIX_PRIORITIZED_IO` and `_POSIX_PRIORITY_SCHEDULING` are defined, then
 1882 asynchronous I/O is queued in priority order, with the priority of each asynchronous operation
 1883 based on the current scheduling priority of the calling process. The *aio_reqprio* member can be
 1884 used to lower (but not raise) the asynchronous I/O operation priority and is within the range
 1885 zero through `{AIO_PRIO_DELTA_MAX}`, inclusive. Unless both `_POSIX_PRIORITIZED_IO` and
 1886 `_POSIX_PRIORITY_SCHEDULING` are defined, the order of processing asynchronous I/O
 1887 requests is unspecified. When both `_POSIX_PRIORITIZED_IO` and
 1888 `_POSIX_PRIORITY_SCHEDULING` are defined, the order of processing of requests submitted
 1889 by processes whose schedulers are not `SCHED_FIFO`, `SCHED_RR`, or `SCHED_SPORADIC` is
 1890 unspecified. The priority of an asynchronous request is computed as (process scheduling
 1891 priority) minus *aio_reqprio*. The priority assigned to each asynchronous I/O request is an
 1892 indication of the desired order of execution of the request relative to other asynchronous I/O
 1893 requests for this file. If `_POSIX_PRIORITIZED_IO` is defined, requests issued with the same
 1894 priority to a character special file are processed by the underlying device in FIFO order; the order
 1895 of processing of requests of the same priority issued to files that are not character special files is
 1896 unspecified. Numerically higher priority values indicate requests of higher priority. The value of
 1897 *aio_reqprio* has no effect on process scheduling priority. When prioritized asynchronous I/O
 1898 requests to the same file are blocked waiting for a resource required for that I/O operation, the
 1899 higher-priority I/O requests shall be granted the resource before lower-priority I/O requests are
 1900 granted the resource. The relative priority of asynchronous I/O and synchronous I/O is
 1901 implementation-defined. If `_POSIX_PRIORITIZED_IO` is defined, the implementation shall
 1902 define for which files I/O prioritization is supported.

1903 The *aio_sigevent* determines how the calling process shall be notified upon I/O completion, as
 1904 specified in Section 2.4.1 (on page 528). If *aio_sigevent.sigev_notify* is `SIGEV_NONE`, then no
 1905 signal shall be posted upon I/O completion, but the error status for the operation and the return
 1906 status for the operation shall be set appropriately.

1907 The *aio_lio_opcode* field is used only by the `lio_listio()` call. The `lio_listio()` call allows multiple
 1908 asynchronous I/O operations to be submitted at a single time. The function takes as an
 1909 argument an array of pointers to **aiocb** structures. Each **aiocb** structure indicates the operation to
 1910 be performed (read or write) via the *aio_lio_opcode* field.

1911 The address of the **aiocb** structure is used as a handle for retrieving the error status and return
 1912 status of the asynchronous operation while it is in progress.

1913 The **aiocb** structure and the data buffers associated with the asynchronous I/O operation are
 1914 being used by the system for asynchronous I/O while, and only while, the error status of the
 1915 asynchronous operation is equal to `EINPROGRESS`. Applications shall not modify the **aiocb**
 1916 structure while the structure is being used by the system for asynchronous I/O.

1917 The return status of the asynchronous operation is the number of bytes transferred by the I/O
 1918 operation. If the error status is set to indicate an error completion, then the return status is set to
 1919 the return value that the corresponding `read()`, `write()`, or `fsync()` call would have returned.
 1920 When the error status is not equal to `EINPROGRESS`, the return status shall reflect the return
 1921 status of the corresponding synchronous operation.

1922 2.8.3 Memory Management

1923 2.8.3.1 Memory Locking

1924 ML The functionality described in this section is dependent on support of the Process Memory
1925 Locking option (and the rest of this section is not further shaded for this option).

1926 Range memory locking operations are defined in terms of pages. Implementations may restrict
1927 the size and alignment of range lockings to be on page-size boundaries. The page size, in bytes,
1928 is the value of the configurable system variable {PAGESIZE}. If an implementation has no
1929 restrictions on size or alignment, it may specify a 1-byte page size.

1930 Memory locking guarantees the residence of portions of the address space. It is
1931 implementation-defined whether locking memory guarantees fixed translation between virtual
1932 addresses (as seen by the process) and physical addresses. Per-process memory locks are not
1933 inherited across a *fork()*, and all memory locks owned by a process are unlocked upon *exec* or
1934 process termination. Unmapping of an address range removes any memory locks established on
1935 that address range by this process.

1936 2.8.3.2 Memory Mapped Files

1937 MF The functionality described in this section is dependent on support of the Memory Mapped Files
1938 option (and the rest of this section is not further shaded for this option).

1939 Range memory mapping operations are defined in terms of pages. Implementations may
1940 restrict the size and alignment of range mappings to be on page-size boundaries. The page size,
1941 in bytes, is the value of the configurable system variable {PAGESIZE}. If an implementation has
1942 no restrictions on size or alignment, it may specify a 1-byte page size.

1943 Memory mapped files provide a mechanism that allows a process to access files by directly
1944 incorporating file data into its address space. Once a file is mapped into a process address space,
1945 the data can be manipulated as memory. If more than one process maps a file, its contents are
1946 shared among them. If the mappings allow shared write access, then data written into the
1947 memory object through the address space of one process appears in the address spaces of all
1948 processes that similarly map the same portion of the memory object.

1949 SHM Shared memory objects are named regions of storage that may be independent of the file system
1950 and can be mapped into the address space of one or more processes to allow them to share the
1951 associated memory.

1952 SHM An *unlink()* of a file or *shm_unlink()* of a shared memory object, while causing the removal of the
1953 name, does not unmap any mappings established for the object. Once the name has been
1954 removed, the contents of the memory object are preserved as long as it is referenced. The
1955 memory object remains referenced as long as a process has the memory object open or has some
1956 area of the memory object mapped.

1957 2.8.3.3 Memory Protection

1958 MPR MF The functionality described in this section is dependent on support of the Memory Protection
1959 and Memory Mapped Files option (and the rest of this section is not further shaded for these
1960 options).

1961 When an object is mapped, various application accesses to the mapped region may result in
1962 signals. In this context, SIGBUS is used to indicate an error using the mapped object, and
1963 SIGSEGV is used to indicate a protection violation or misuse of an address:

- 1964 • A mapping may be restricted to disallow some types of access.

- 1965 • Write attempts to memory that was mapped without write access, or any access to memory
1966 mapped PROT_NONE, shall result in a SIGSEGV signal.
- 1967 • References to unmapped addresses shall result in a SIGSEGV signal.
- 1968 • Reference to whole pages within the mapping, but beyond the current length of the object,
1969 shall result in a SIGBUS signal.
- 1970 • The size of the object is unaffected by access beyond the end of the object (even if a SIGBUS is
1971 not generated).

1972 2.8.3.4 *Typed Memory Objects*

1973 TYM The functionality described in this section is dependent on support of the Typed Memory
1974 Objects option (and the rest of this section is not further shaded for this option).

1975 Implementations may support the Typed Memory Objects option without supporting the
1976 Memory Mapped Files option or the Shared Memory Objects option. Typed memory objects are
1977 implementation-configurable named storage pools accessible from one or more processors in a
1978 system, each via one or more ports, such as backplane buses, LANs, I/O channels, and so on.
1979 Each valid combination of a storage pool and a port is identified through a name that is defined
1980 at system configuration time, in an implementation-defined manner; the name may be
1981 independent of the file system. Using this name, a typed memory object can be opened and
1982 mapped into process address space. For a given storage pool and port, it is necessary to support
1983 both dynamic allocation from the pool as well as mapping at an application-supplied offset
1984 within the pool; when dynamic allocation has been performed, subsequent deallocation must be
1985 supported. Lastly, accessing typed memory objects from different ports requires a method for
1986 obtaining the offset and length of contiguous storage of a region of typed memory (dynamically
1987 allocated or not); this allows typed memory to be shared among processes and/or processors
1988 while being accessed from the desired port.

1989 2.8.4 **Process Scheduling**

1990 **Scheduling Policies**

1991 PS The functionality described in this section is dependent on support of the Process Scheduling
1992 option (and the rest of this section is not further shaded for this option).

1993 The scheduling semantics described in this volume of IEEE Std. 1003.1-200x are defined in terms
1994 of a conceptual model that contains a set of thread lists. No implementation structures are
1995 necessarily implied by the use of this conceptual model. It is assumed that no time elapses
1996 during operations described using this model, and therefore no simultaneous operations are
1997 possible. This model discusses only processor scheduling for runnable threads, but it should be
1998 noted that greatly enhanced predictability of realtime applications result if the sequencing of
1999 other resources takes processor scheduling policy into account.

2000 There is, conceptually, one thread list for each priority. Any runnable thread may be on any
2001 thread list. Multiple scheduling policies shall be provided. Each non-empty thread list is
2002 ordered, contains a head as one end of its order, and a tail as the other. The purpose of a
2003 scheduling policy is to define the allowable operations on this set of lists (for example, moving
2004 threads between and within lists).

2005 Each process shall be controlled by an associated scheduling policy and priority. These
2006 parameters may be specified by explicit application execution of the *sched_setscheduler()* or
2007 *sched_setparam()* functions.

2008 Each thread shall be controlled by an associated scheduling policy and priority. These
 2009 parameters may be specified by explicit application execution of the *pthread_setschedparam()*
 2010 function.

2011 Associated with each policy is a priority range. Each policy definition shall specify the minimum
 2012 priority range for that policy. The priority ranges for each policy may but need not overlap the
 2013 priority ranges of other policies.

2014 A conforming implementation shall select the thread that is defined as being at the head of the
 2015 highest priority non-empty thread list to become a running thread, regardless of its associated
 2016 policy. This thread is then removed from its thread list.

2017 Four scheduling policies are specifically required. Other implementation-defined scheduling
 2018 policies may be defined. The following symbols are defined in the Base Definitions volume of
 2019 IEEE Std. 1003.1-200x, <**sched.h**>:

2020 **SCHED_FIFO** First in, first out (FIFO) scheduling policy.

2021 **SCHED_RR** Round robin scheduling policy.

2022 ss **SCHED_SPORADIC** Sporadic server scheduling policy.

2023 **SCHED_OTHER** Another scheduling policy.

2024 The values of these symbols shall be distinct.

2025 **SCHED_FIFO**

2026 Conforming implementations shall include a scheduling policy called the FIFO scheduling
 2027 policy.

2028 Threads scheduled under this policy are chosen from a thread list that is ordered by the time its
 2029 threads have been on the list without being executed; generally, the head of the list is the thread
 2030 that has been on the list the longest time, and the tail is the thread that has been on the list the
 2031 shortest time.

2032 Under the **SCHED_FIFO** policy, the modification of the definitional thread lists is as follows:

- 2033 1. When a running thread becomes a preempted thread, it becomes the head of the thread list
 2034 for its priority.
- 2035 2. When a blocked thread becomes a runnable thread, it becomes the tail of the thread list for
 2036 its priority.
- 2037 3. When a running thread calls the *sched_setscheduler()* function, the process specified in the
 2038 function call is modified to the specified policy and the priority specified by the *param*
 2039 argument.
- 2040 4. When a running thread calls the *sched_setparam()* function, the priority of the process
 2041 specified in the function call is modified to the priority specified by the *param* argument.
- 2042 5. When a running thread calls the *pthread_setschedparam()* function, the thread specified in
 2043 the function call is modified to the specified policy and the priority specified by the *param*
 2044 argument.
- 2045 6. If a thread whose policy or priority has been modified is a running thread or is runnable, it
 2046 then becomes the tail of the thread list for its new priority.
- 2047 7. When a running thread issues the *sched_yield()* function, the thread becomes the tail of the
 2048 thread list for its priority.

2049 8. At no other time is the position of a thread with this scheduling policy within the thread
2050 lists affected.

2051 For this policy, valid priorities shall be within the range returned by the *sched_get_priority_max()*
2052 and *sched_get_priority_min()* functions when SCHED_FIFO is provided as the parameter.
2053 Conforming implementations shall provide a priority range of at least 32 priorities for this
2054 policy.

2055 SCHED_RR

2056 Conforming implementations shall include a scheduling policy called the *round robin* scheduling
2057 policy. This policy is identical to the SCHED_FIFO policy with the additional condition that
2058 when the implementation detects that a running thread has been executing as a running thread
2059 for a time period of the length returned by the *sched_rr_get_interval()* function or longer, the
2060 thread shall become the tail of its thread list and the head of that thread list shall be removed
2061 and made a running thread.

2062 The effect of this policy is to ensure that if there are multiple SCHED_RR threads at the same
2063 priority, one of them does not monopolize the processor. An application should not rely only on
2064 the use of SCHED_RR to ensure application progress among multiple threads if the application
2065 includes threads using the SCHED_FIFO policy at the same or higher priority levels or
2066 SCHED_RR threads at a higher priority level.

2067 A thread under this policy that is preempted and subsequently resumes execution as a running
2068 thread completes the unexpired portion of its round robin interval time period.

2069 For this policy, valid priorities shall be within the range returned by the *sched_get_priority_max()*
2070 and *sched_get_priority_min()* functions when SCHED_RR is provided as the parameter.
2071 Conforming implementations shall provide a priority range of at least 32 priorities for this
2072 policy.

2073 SCHED_SPORADIC

2074 SS|TSP The functionality described in this section is dependent on support of the Process Sporadic
2075 Server or Thread Sporadic Server options (and the rest of this section is not further shaded for
2076 these options).

2077 If `_POSIX_SPORADIC_SERVER` or `_POSIX_THREAD_SPORADIC_SERVER` is defined, the
2078 implementation shall include a scheduling policy identified by the value `SCHED_SPORADIC`.

2079 The sporadic server policy is based primarily on two parameters: the *replenishment period* and the
2080 *available execution capacity*. The replenishment period is given by the *sched_ss_repl_period*
2081 member of the **sched_param** structure. The available execution capacity is initialized to the
2082 value given by the *sched_ss_init_budget* member of the same parameter. The sporadic server
2083 policy is identical to the SCHED_FIFO policy with some additional conditions that cause the
2084 thread's assigned priority to be switched between the values specified by the *sched_priority* and
2085 *sched_ss_low_priority* members of the **sched_param** structure.

2086 The priority assigned to a thread using the sporadic server scheduling policy is determined in
2087 the following manner: if the available execution capacity is greater than zero and the number of
2088 pending replenishment operations is strictly less than *sched_ss_max_repl*, the thread is assigned
2089 the priority specified by *sched_priority*; otherwise, the assigned priority shall be
2090 *sched_ss_low_priority*. If the value of *sched_priority* is less than or equal to the value of
2091 *sched_ss_low_priority*, the results are undefined. When active, the thread shall belong to the
2092 thread list corresponding to its assigned priority level, according to the mentioned priority
2093 assignment. The modification of the available execution capacity and, consequently of the
2094 assigned priority, is done as follows:

- 2095 1. When the thread at the head of the *sched_priority* list becomes a running thread, its
 2096 execution time shall be limited to at most its available execution capacity, plus the
 2097 resolution of the execution time clock used for this scheduling policy. This resolution shall
 2098 be implementation-defined.
- 2099 2. Each time the thread is inserted at the tail of the list associated with *sched_priority*—
 2100 because as a blocked thread it became runnable with priority *sched_priority* or because a
 2101 replenishment operation was performed—the time at which this operation is done is
 2102 posted as the *activation_time*.
- 2103 3. When the running thread with assigned priority equal to *sched_priority* becomes a
 2104 preempted thread, it becomes the head of the thread list for its priority, and the execution
 2105 time consumed is subtracted from the available execution capacity. If the available
 2106 execution capacity would become negative by this operation, it shall be set to zero.
- 2107 4. When the running thread with assigned priority equal to *sched_priority* becomes a blocked
 2108 thread, the execution time consumed is subtracted from the available execution capacity,
 2109 and a replenishment operation is scheduled, as described in 6 and 7. If the available
 2110 execution capacity would become negative by this operation, it shall be set to zero.
- 2111 5. When the running thread with assigned priority equal to *sched_priority* reaches the limit
 2112 imposed on its execution time, it becomes the tail of the thread list for
 2113 *sched_ss_low_priority*, the execution time consumed is subtracted from the available
 2114 execution capacity (which becomes zero), and a replenishment operation is scheduled, as
 2115 described in 6 and 7.
- 2116 6. Each time a replenishment operation is scheduled, the amount of execution capacity to be
 2117 replenished, *replenish_amount*, is set equal to the execution time consumed by the thread
 2118 since the *activation_time*. The replenishment is scheduled to occur at *activation_time* plus
 2119 *sched_ss_repl_period*. If the scheduled time obtained is before the current time, the
 2120 replenishment operation is carried out immediately. Several replenishment operations may
 2121 be pending at the same time, each of which will be serviced at its respective scheduled
 2122 time. With the above rules, the number of replenishment operations simultaneously
 2123 pending for a given thread that is scheduled under the sporadic server policy shall not be
 2124 greater than *sched_ss_max_repl*.
- 2125 7. A replenishment operation consists of adding the corresponding *replenish_amount* to the
 2126 available execution capacity at the scheduled time. If, as a consequence of this operation,
 2127 the execution capacity would become larger than *sched_ss_initial_budget*, it shall be
 2128 rounded down to a value equal to *sched_ss_initial_budget*. Additionally, if the thread was
 2129 runnable or running, and had assigned priority equal to *sched_ss_low_priority*, then it
 2130 becomes the tail of the thread list for *sched_priority*.
- 2131 Execution time is defined in Section 2.2.2 (on page 513).
- 2132 For this policy, changing the value of a CPU-time clock via *clock_settime()* shall have no effect on
 2133 its behavior.
- 2134 For this policy, valid priorities shall be within the range returned by the *sched_get_priority_min()*
 2135 and *sched_get_priority_max()* functions when SCHED_SPORADIC is provided as the parameter.
 2136 Conforming implementations shall provide a priority range of at least 32 distinct priorities for
 2137 this policy.

2138 **SCHED_OTHER**

2139 Conforming implementations shall include one scheduling policy identified as SCHED_OTHER
 2140 (which may execute identically with either the FIFO or round robin scheduling policy). The
 2141 effect of scheduling threads with the SCHED_OTHER policy in a system in which other threads
 2142 ss are executing under SCHED_FIFO, SCHED_RR, or SCHED_SPORADIC is implementation-
 2143 defined.

2144 This policy is defined to allow conforming applications to be able to indicate that they no longer
 2145 need a realtime scheduling policy in a portable manner.

2146 For threads executing under this policy, the implementation shall use only priorities within the
 2147 range returned by the *sched_get_priority_max()* and *sched_get_priority_min()* functions when
 2148 SCHED_OTHER is provided as the parameter.

2149 **2.8.5 Clocks and Timers**

2150 TMR The functionality described in this section is dependent on support of the Timers option (and the
 2151 rest of this section is not further shaded for this option).

2152 The <time.h> header defines the types and manifest constants used by the timing facility.

2153 **Time Value Specification Structures**

2154 Many of the timing facility functions accept or return time value specifications. A time value
 2155 structure **timespec** specifies a single time value and includes at least the following members:

2156

2157

2158

2159

Member Type	Member Name	Description
time_t	<i>tv_sec</i>	Seconds.
long	<i>tv_nsec</i>	Nanoseconds.

2160 The *tv_nsec* member is only valid if greater than or equal to zero, and less than the number of
 2161 nanoseconds in a second (1,000 million). The time interval described by this structure is (*tv_sec* *
 2162 10⁹ + *tv_nsec*) nanoseconds.

2163 A time value structure **itimerspec** specifies an initial timer value and a repetition interval for use
 2164 by the per-process timer functions. This structure includes at least the following members:

2165

2166

2167

2168

Member Type	Member Name	Description
struct timespec	<i>it_interval</i>	Timer period.
struct timespec	<i>it_value</i>	Timer expiration.

2169 If the value described by *it_value* is non-zero, it indicates the time to or time of the next timer
 2170 expiration (for relative and absolute timer values, respectively). If the value described by *it_value*
 2171 is zero, the timer shall be disarmed.

2172 If the value described by *it_interval* is non-zero, it specifies an interval to be used in reloading the
 2173 timer when it expires; that is, a periodic timer is specified. If the value described by *it_interval* is
 2174 zero, the timer is disarmed after its next expiration; that is, a one-shot timer is specified.

2175 **Timer Event Notification Control Block**

2176 RTS Per-process timers may be created that notify the process of timer expirations by queuing a
 2177 realtime extended signal. The **sigevent** structure, defined in the Base Definitions volume of
 2178 IEEE Std. 1003.1-200x, <**signal.h**>, is used in creating such a timer. The **sigevent** structure
 2179 contains the signal number and an application-specific data value to be used when notifying the
 2180 calling process of timer expiration events.

2181 **Manifest Constants**

2182 The following constants are defined in the Base Definitions volume of IEEE Std. 1003.1-200x,
 2183 <**time.h**>:

2184 **CLOCK_REALTIME** The identifier for the system-wide realtime clock.

2185 **TIMER_ABSTIME** Flag indicating time is absolute with respect to the clock associated
 2186 with a timer.

2187 MON **CLOCK_MONOTONIC** The identifier for the system-wide monotonic clock, which is defined
 2188 as a clock whose value cannot be set via *clock_settime()* and which
 2189 cannot have backward clock jumps. The maximum possible clock
 2190 jump is implementation-defined.

2191 MON The maximum allowable resolution for the **CLOCK_REALTIME** and the
 2192 **CLOCK_MONOTONIC** clocks and all time services based on these clocks is represented by
 2193 {**_POSIX_CLOCKRES_MIN**} and is defined as 20ms (1/50 of a second). Implementations may
 2194 support smaller values of resolution for these clocks to provide finer granularity time bases. The
 2195 actual resolution supported by an implementation for a specific clock is obtained using the
 2196 *clock_getres()* function. If the actual resolution supported for a time service based on one of these
 2197 clocks differs from the resolution supported for that clock, the implementation shall document
 2198 this difference.

2199 MON The minimum allowable maximum value for the **CLOCK_REALTIME** and the
 2200 **CLOCK_MONOTONIC** clocks and all absolute time services based on them is the same as that
 2201 defined by the ISO C standard for the **time_t** type. If the maximum value supported by a time
 2202 service based on one of these clocks differs from the maximum value supported by that clock,
 2203 the implementation shall document this difference.

2204 **Execution Time Monitoring**

2205 CPT If **_POSIX_CPUTIME** is defined, process CPU-time clocks shall be supported in addition to the
 2206 clocks described in **Manifest Constants**.

2207 TCT If **_POSIX_THREAD_CPUTIME** is defined, thread CPU-time clocks shall be supported.

2208 CPT|TCT CPU-time clocks measure execution or CPU time, which is defined in the Base Definitions
 2209 volume of IEEE Std. 1003.1-200x, Section 3.120, CPU Time (Execution Time). The mechanism
 2210 used to measure execution time is described in the Base Definitions volume of
 2211 IEEE Std. 1003.1-200x, Section 4.7, Measurement of Execution Time.

2212 CPT If **_POSIX_CPUTIME** is defined, the following constant of the type **clockid_t** shall be defined in
 2213 <**time.h**>:

2214 **CLOCK_PROCESS_CPUTIME_ID**

2215 When this value of the type **clockid_t** is used in a *clock()* or *timer*()* function call, it is
 2216 interpreted as the identifier of the CPU-time clock associated with the process making the
 2217 function call.
 2218

2219 TCT If `_POSIX_THREAD_CPUTIME` is defined, the following constant of the type `clockid_t` shall be
2220 defined in `<time.h>`:

2221 `CLOCK_THREAD_CPUTIME_ID`

2222 When this value of the type `clockid_t` is used in a `clock()` or `timer*()` function call, it is
2223 interpreted as the identifier of the CPU-time clock associated with the thread making the
2224 function call.
2225

2226 **2.9 Threads**

2227 THR The functionality described in this section is dependent on support of the Threads option (and
2228 the rest of this section is not further shaded for this option).

2229 This section defines functionality to support multiple flows of control, called *threads*, within a
2230 process. For the definition of *threads*, see the Base Definitions volume of IEEE Std. 1003.1-200x,
2231 Section 3.395, Thread.

2232 The specific functional areas covered by threads and their scope includes:

- 2233 • Thread management: the creation, control, and termination of multiple flows of control in the
2234 same process under the assumption of a common shared address space
- 2235 • Synchronization primitives optimized for tightly coupled operation of multiple control flows
2236 in a common, shared address space

2237 **2.9.1 Thread-Safety**

2238 All functions defined by this volume of IEEE Std. 1003.1-200x shall be thread-safe, except that
2239 the following functions need not be thread-safe.

2240	<i>asctime()</i>	<i>gethostbyname()</i>	<i>getprotobynumber()</i>	<i>inet_ntoa()</i>	<i>ttyname()</i>
2241	<i>ctime()</i>	<i>gethostent()</i>	<i>getprotoent()</i>	<i>localtime()</i>	<i>unsetenv()</i>
2242	<i>getc_unlocked()</i>	<i>getlogin()</i>	<i>getpwnam()</i>	<i>putc_unlocked()</i>	<i>wcstombs()</i>
2243	<i>getchar_unlocked()</i>	<i>getnetbyaddr()</i>	<i>getpwuid()</i>	<i>putchar_unlocked()</i>	<i>wctomb()</i>
2244	<i>getenv()</i>	<i>getnetbyname()</i>	<i>getservbyname()</i>	<i>rand()</i>	
2245	<i>getgrgid()</i>	<i>getnetent()</i>	<i>getservbyport()</i>	<i>readdir()</i>	
2246	<i>getgrnam()</i>	<i>getopt()</i>	<i>getservent()</i>	<i>setenv()</i>	
2247	<i>gethostbyaddr()</i>	<i>getprotobyname()</i>	<i>gmtime()</i>	<i>strtok()</i>	

2248 XSI	<i>basename()</i>	<i>dbm_open()</i>	<i>fcvt()</i>	<i>hdestroy()</i>	<i>setgrent()</i>
2249	<i>catgets()</i>	<i>dbm_store()</i>	<i>gcvt()</i>	<i>hsearch()</i>	<i>setkey()</i>
2250	<i>crypt()</i>	<i>dirname()</i>	<i>getdate()</i>	<i>l64a()</i>	<i>setpwent()</i>
2251	<i>dbm_clearerr()</i>	<i>derror()</i>	<i>getenv()</i>	<i>lgamma()</i>	<i>setutxent()</i>
2252	<i>dbm_close()</i>	<i>drand48()</i>	<i>getgrent()</i>	<i>lrand48()</i>	<i>strerror()</i>
2253	<i>dbm_delete()</i>	<i>ecvt()</i>	<i>getpwent()</i>	<i>mrnd48()</i>	
2254	<i>dbm_error()</i>	<i>encrypt()</i>	<i>getutxent()</i>	<i>nl_langinfo()</i>	
2255	<i>dbm_fetch()</i>	<i>endgrent()</i>	<i>getutxid()</i>	<i>ptsname()</i>	
2256	<i>dbm_firstkey()</i>	<i>endpwent()</i>	<i>getutxline()</i>	<i>putenv()</i>	
2257	<i>dbm_nextkey()</i>	<i>endutxent()</i>	<i>hcreate()</i>	<i>pututxline()</i>	

2258

2259 The *read()* function need not be thread-safe when reading from a pipe, FIFO, socket, or terminal
2260 device.

2261 **Note:** While a read from a pipe of {PIPE_MAX}*2 bytes may not generate a single atomic
2262 and thread-safe stream of bytes, it should generate “several” (individually atomic)
2263 thread-safe streams of bytes. Similarly, while reading from a terminal device may
2264 not generate a single atomic and thread-safe stream of bytes, it should generate some
2265 finite number of (individually atomic) and thread-safe streams of bytes. That is,
2266 concurrent calls to read for a pipe, FIFO, or terminal device are not allowed to result
2267 in corrupting the stream of bytes or other internal data. However, *read()*, in these
2268 cases, is not required to return a single contiguous and atomic stream of bytes.

2269 The *ctermid()* and *tmpnam()* functions need not be thread-safe if passed a NULL argument. The
 2270 *wcrtomb()* and *wcsrtombs()* functions need not be thread-safe if passed a NULL *ps* argument.

2271 Implementations shall provide internal synchronization as necessary in order to satisfy this
 2272 requirement.

2273 2.9.2 Thread IDs

2274 Although implementations may have thread IDs that are unique in a system, applications
 2275 should only assume that thread IDs are usable and unique within a single process. The effect of
 2276 calling any of the functions defined in this volume of IEEE Std. 1003.1-200x and passing as an
 2277 argument the thread ID of a thread from another process is unspecified. A conforming
 2278 implementation is free to reuse a thread ID after the thread terminates if it was created with the
 2279 *detachstate* attribute set to `PTHREAD_CREATE_DETACHED` or if *pthread_detach()* or
 2280 *pthread_join()* has been called for that thread. If a thread is detached, its thread ID is invalid for
 2281 use as an argument in a call to *pthread_detach()* or *pthread_join()*.

2282 2.9.3 Thread Mutexes

2283 A thread that has blocked shall not prevent any unblocked thread that is eligible to use the same
 2284 processing resources from eventually making forward progress in its execution. Eligibility for
 2285 processing resources is determined by the scheduling policy.

2286 A thread becomes the owner of a mutex, *m*, when one of the following occurs:

- 2287 • It returns successfully from *pthread_mutex_lock()* with *m* as the *mutex* argument.
- 2288 • It returns successfully from *pthread_mutex_trylock()* with *m* as the *mutex* argument.
- 2289 TMO • It returns successfully from *pthread_mutex_timedwait()* with *m* as the *mutex* argument.
- 2290 • It returns (successfully or not) from *pthread_cond_wait()* with *m* as the *mutex* argument
 2291 (except as explicitly indicated otherwise for certain errors).
- 2292 • It returns (successfully or not) from *pthread_cond_timedwait()* with *m* as the *mutex* argument
 2293 (except as explicitly indicated otherwise for certain errors).

2294 The thread remains the owner of *m* until one of the following occurs:

- 2295 • It executes *pthread_mutex_unlock()* with *m* as the *mutex* argument
- 2296 • It blocks in a call to *pthread_cond_wait()* with *m* as the *mutex* argument.
- 2297 • It blocks in a call to *pthread_cond_timedwait()* with *m* as the *mutex* argument.

2298 The implementation behaves as if at all times there is at most one owner of any mutex.

2299 A thread that becomes the owner of a mutex is said to have *acquired* the mutex and the mutex is
 2300 said to have become *locked*; when a thread gives up ownership of a mutex it is said to have
 2301 *released* the mutex and the mutex is said to have become *unlocked*.

2302 **2.9.4 Thread Scheduling**2303 **Thread Scheduling Attributes**

2304 Thread scheduling attributes are dependent on support of the Thread Execution Scheduling
2305 option.

2306 In support of the scheduling function, threads have attributes which are accessed through the
2307 *pthread_attr_t* thread creation attributes object.

2308 The *contentionscope* attribute defines the scheduling contention scope of the thread to be either
2309 PTHREAD_SCOPE_PROCESS or PTHREAD_SCOPE_SYSTEM.

2310 The *inheritsched* attribute specifies whether a newly created thread is to inherit the scheduling
2311 attributes of the creating thread or to have its scheduling values set according to the other
2312 scheduling attributes in the *pthread_attr_t* object.

2313 The *schedpolicy* attribute defines the scheduling policy for the thread. The *schedparam* attribute
2314 defines the scheduling parameters for the thread. The interaction of threads having different
2315 policies within a process is described as part of the definition of those policies.

2316 If the Thread Execution Scheduling option is defined, and the *schedpolicy* attribute specifies one
2317 of the priority-based policies defined under this option, the *schedparam* attribute contains the
2318 scheduling priority of the thread. A conforming implementation ensures that the priority value
2319 in *schedparam* is in the range associated with the scheduling policy when the thread attributes
2320 object is used to create a thread, or when the scheduling attributes of a thread are dynamically
2321 modified. The meaning of the priority value in *schedparam* is the same as that of *priority*.

2322 TSP If *_POSIX_THREAD_SPORADIC_SERVER* is defined, the *schedparam* attribute supports four
2323 new members that are used for the sporadic server scheduling policy. These members are
2324 *sched_ss_low_priority*, *sched_ss_repl_period*, *sched_ss_init_budget*, and *sched_ss_max_repl*. The
2325 meaning of these attributes is the same as in the definitions that appear under Section 2.8.4 (on
2326 page 546).

2327 When a process is created, its single thread has a scheduling policy and associated attributes
2328 equal to the process' policy and attributes. The default scheduling contention scope value is
2329 implementation-defined. The default values of other scheduling attributes are implementation-
2330 defined.

2331 **Thread Scheduling Contention Scope**

2332 The scheduling contention scope of a thread defines the set of threads with which the thread
2333 competes for use of the processing resources. The scheduling operation selects at most one
2334 thread to execute on each processor at any point in time and the thread's scheduling attributes
2335 (for example, *priority*), whether under process scheduling contention scope or system scheduling
2336 contention scope, are the parameters used to determine the scheduling decision.

2337 The scheduling contention scope, in the context of scheduling a mixed scope environment,
2338 effects threads as follows:

- 2339 • A thread created with PTHREAD_SCOPE_SYSTEM scheduling contention scope contends
2340 for resources with all other threads in the same scheduling allocation domain relative to their
2341 system scheduling attributes. The system scheduling attributes of a thread created with
2342 PTHREAD_SCOPE_SYSTEM scheduling contention scope are the scheduling attributes with
2343 which the thread was created. The system scheduling attributes of a thread created with
2344 PTHREAD_SCOPE_PROCESS scheduling contention scope are the implementation-defined
2345 mapping into system attribute space of the scheduling attributes with which the thread was

- 2346 created.
- 2347 • Threads created with PTHREAD_SCOPE_PROCESS scheduling contention scope contend
2348 directly with other threads within their process that were created with
2349 PTHREAD_SCOPE_PROCESS scheduling contention scope. The contention is resolved
2350 based on the threads' scheduling attributes and policies. It is unspecified how such threads
2351 are scheduled relative to threads in other processes or threads with
2352 PTHREAD_SCOPE_SYSTEM scheduling contention scope.
- 2353 • Conforming implementations shall support the PTHREAD_SCOPE_PROCESS scheduling
2354 contention scope, the PTHREAD_SCOPE_SYSTEM scheduling contention scope, or both.

2355 Scheduling Allocation Domain

2356 Implementations shall support scheduling allocation domains containing one or more
2357 processors. It should be noted that the presence of multiple processors does not automatically
2358 indicate a scheduling allocation domain size greater than one. Conforming implementations on
2359 multiprocessors may map all or any subset of the CPUs to one or multiple scheduling allocation
2360 domains, and could define these scheduling allocation domains on a per-thread, per-process, or
2361 per-system basis, depending on the types of applications intended to be supported by the
2362 implementation. The scheduling allocation domain is independent of scheduling contention
2363 scope, as the scheduling contention scope merely defines the set of threads with which a thread
2364 contends for processor resources, while scheduling allocation domain defines the set of
2365 processors for which it contends. The semantics of how this contention is resolved among
2366 threads for processors is determined by the scheduling policies of the threads.

2367 The choice of scheduling allocation domain size and the level of application control over
2368 scheduling allocation domains is implementation-defined. Conforming implementations may
2369 change the size of scheduling allocation domains and the binding of threads to scheduling
2370 allocation domains at any time.

2371 For application threads with scheduling allocation domains of size equal to one, the scheduling
2372 rules defined for SCHED_FIFO and SCHED_RR shall be used; see **Scheduling Policies** (on page
2373 546). All threads with system scheduling contention scope, regardless of the processes in which
2374 they reside, compete for the processor according to their priorities. Threads with process
2375 scheduling contention scope compete only with other threads with process scheduling
2376 contention scope within their process.

2377 For application threads with scheduling allocation domains of size greater than one, the rules
2378 TSP defined for SCHED_FIFO, SCHED_RR, and SCHED_SPORADIC shall be used in an
2379 implementation-defined manner. Each thread with system scheduling contention scope
2380 competes for the processors in its scheduling allocation domain in an implementation-defined
2381 manner according to its priority. Threads with process scheduling contention scope are
2382 scheduled relative to other threads within the same scheduling contention scope in the process.

2383 TSP If _POSIX_THREAD_SPORADIC_SERVER is defined, the rules defined for SCHED_SPORADIC
2384 in **Scheduling Policies** (on page 546) shall be used in an implementation-defined manner for
2385 application threads whose scheduling allocation domain size is greater than one.

2386 **Scheduling Documentation**

2387 If `_POSIX_PRIORITY_SCHEDULING` is defined, then any scheduling policies beyond
 2388 TSP `SCHED_OTHER`, `SCHED_FIFO`, `SCHED_RR`, and `SCHED_SPORADIC`, as well as the effects of
 2389 the scheduling policies indicated by these other values, and the attributes required in order to
 2390 support such a policy, are implementation-defined. Furthermore, the implementation shall
 2391 document the effect of all processor scheduling allocation domain values supported for these
 2392 policies.

2393 **2.9.5 Thread Cancellation**

2394 The thread cancellation mechanism allows a thread to terminate the execution of any other
 2395 thread in the process in a controlled manner. The target thread (that is, the one that is being
 2396 canceled) is allowed to hold cancellation requests pending in a number of ways and to perform
 2397 application-specific cleanup processing when the notice of cancellation is acted upon.

2398 Cancellation is controlled by the cancellation control functions. Each thread maintains its own
 2399 cancelability state. Cancellation may only occur at cancellation points or when the thread is
 2400 asynchronously cancelable.

2401 The thread cancellation mechanism described in this section depends upon programs having set
 2402 *deferred cancelability* state, which is specified as the default. Applications shall also carefully
 2403 follow static lexical scoping rules in their execution behavior. For example, use of *setjmp()*,
 2404 *return*, *goto*, and so on, to leave user-defined cancellation scopes without doing the necessary
 2405 scope pop operation results in undefined behavior.

2406 Use of asynchronous cancelability while holding resources which potentially need to be released
 2407 may result in resource loss. Similarly, cancellation scopes may only be safely manipulated
 2408 (pushed and popped) when the thread is in the *deferred* or *disabled* cancelability states.

2409 **2.9.5.1 Cancelability States**

2410 The cancelability state of a thread determines the action taken upon receipt of a cancellation
 2411 request. The thread may control cancellation in a number of ways.

2412 Each thread maintains its own cancelability state, which may be encoded in two bits:

- 2413 1. **Cancelability-Enable:** When cancelability is `PTHREAD_CANCEL_DISABLE` (as defined in
 2414 the Base Definitions volume of IEEE Std. 1003.1-200x, `<pthread.h>`), cancellation requests
 2415 against the target thread are held pending. By default, cancelability is set to
 2416 `PTHREAD_CANCEL_ENABLE` (as defined in `<pthread.h>`).
- 2417 2. **Cancelability Type:** When cancelability is enabled and the cancelability type is
 2418 `PTHREAD_CANCEL_ASYNCHRONOUS` (as defined in `<pthread.h>`), new or pending
 2419 cancellation requests may be acted upon at any time. When cancelability is enabled and the
 2420 cancelability type is `PTHREAD_CANCEL_DEFERRED` (as defined in `<pthread.h>`),
 2421 cancellation requests are held pending until a cancellation point (see below) is reached. If
 2422 cancelability is disabled, the setting of the cancelability type has no immediate effect as all
 2423 cancellation requests are held pending; however, once cancelability is enabled again the
 2424 new type is in effect. The cancelability type is `PTHREAD_CANCEL_DEFERRED` in all
 2425 newly created threads including the thread in which *main()* was first invoked.

2426 2.9.5.2 Cancellation Points

2427 Cancellation points occur when a thread is executing the following functions:

2428	<i>accept()</i>	<i>mq_timedsend()</i>	<i>putpmsg()</i>	<i>sigsuspend()</i>
2429	<i>aio_suspend()</i>	<i>msgrcv()</i>	<i>pwrite()</i>	<i>sigtimedwait()</i>
2430	<i>clock_nanosleep()</i>	<i>msgsnd()</i>	<i>read()</i>	<i>sigwait()</i>
2431	<i>close()</i>	<i>msync()</i>	<i>readv()</i>	<i>sigwaitinfo()</i>
2432	<i>connect()</i>	<i>nanosleep()</i>	<i>recv()</i>	<i>sleep()</i>
2433	<i>creat()</i>	<i>open()</i>	<i>recvfrom()</i>	<i>system()</i>
2434	<i>fcntl()</i> ¹	<i>pause()</i>	<i>recvmsg()</i>	<i>tcdrain()</i>
2435	<i>fsync()</i>	<i>poll()</i>	<i>select()</i>	<i>usleep()</i>
2436	<i>getmsg()</i>	<i>pread()</i>	<i>sem_timedwait()</i>	<i>wait()</i>
2437	<i>getpmsg()</i>	<i>pthread_cond_timedwait()</i>	<i>sem_wait()</i>	<i>waitid()</i>
2438	<i>lockf()</i>	<i>pthread_cond_wait()</i>	<i>send()</i>	<i>waitpid()</i>
2439	<i>mq_receive()</i>	<i>pthread_join()</i>	<i>sendmsg()</i>	<i>write()</i>
2440	<i>mq_send()</i>	<i>pthread_testcancel()</i>	<i>sendto()</i>	<i>writev()</i>
2441	<i>mq_timedreceive()</i>	<i>putmsg()</i>	<i>sigpause()</i>	

2442 _____

2443 1. When the *cmd* argument is F_SETLKW.

2444 A cancelation point may also occur when a thread is executing the following functions:

2445	<i>catclose()</i>	<i>ftell()</i>	<i>getwc()</i>	<i>pthread_rwlock_wrlock()</i>
2446	<i>catgets()</i>	<i>ftello()</i>	<i>getwchar()</i>	<i>putc()</i>
2447	<i>catopen()</i>	<i>ftw()</i>	<i>getwd()</i>	<i>putc_unlocked()</i>
2448	<i>closedir()</i>	<i>fwprintf()</i>	<i>glob()</i>	<i>putchar()</i>
2449	<i>closelog()</i>	<i>fwrite()</i>	<i>iconv_close()</i>	<i>putchar_unlocked()</i>
2450	<i>ctermid()</i>	<i>fwscanf()</i>	<i>iconv_open()</i>	<i>puts()</i>
2451	<i>dbm_close()</i>	<i>getc()</i>	<i>ioctl()</i>	<i>pututxline()</i>
2452	<i>dbm_delete()</i>	<i>getc_unlocked()</i>	<i>lseek()</i>	<i>putwc()</i>
2453	<i>dbm_fetch()</i>	<i>getchar()</i>	<i>mkstemp()</i>	<i>putwchar()</i>
2454	<i>dbm_nextkey()</i>	<i>getchar_unlocked()</i>	<i>nftw()</i>	<i>readdir()</i>
2455	<i>dbm_open()</i>	<i>getcwd()</i>	<i>opendir()</i>	<i>readdir_r()</i>
2456	<i>dbm_store()</i>	<i>getdate()</i>	<i>openlog()</i>	<i>remove()</i>
2457	<i>dlclose()</i>	<i>getgrent()</i>	<i>pclose()</i>	<i>rename()</i>
2458	<i>dlopen()</i>	<i>getgrgid()</i>	<i>perror()</i>	<i>rewind()</i>
2459	<i>endgrent()</i>	<i>getgrgid_r()</i>	<i>popen()</i>	<i>rewinddir()</i>
2460	<i>endhostent()</i>	<i>getgrnam()</i>	<i>posix_fadvise()</i>	<i>scanf()</i>
2461	<i>endnetent()</i>	<i>getgrnam_r()</i>	<i>posix_fallocate()</i>	<i>seekdir()</i>
2462	<i>endprotoent()</i>	<i>gethostbyaddr()</i>	<i>posix_madvise()</i>	<i>semop()</i>
2463	<i>endpwent()</i>	<i>gethostbyname()</i>	<i>posix_spawn()</i>	<i>setgrent()</i>
2464	<i>endservent()</i>	<i>gethostent()</i>	<i>posix_spawnnp()</i>	<i>sethostent()</i>
2465	<i>endutxent()</i>	<i>gethostname()</i>	<i>posix_trace_clear()</i>	<i>setnetent()</i>
2466	<i>fclose()</i>	<i>getlogin()</i>	<i>posix_trace_close()</i>	<i>setprotoent()</i>
2467	<i>fcntl()</i> ²	<i>getlogin_r()</i>	<i>posix_trace_create()</i>	<i>setpwent()</i>
2468	<i>fflush()</i>	<i>getnetbyaddr()</i>	<i>posix_trace_create_withlog()</i>	<i>setservent()</i>
2469	<i>fgetc()</i>	<i>getnetbyname()</i>	<i>posix_trace_eventtypelist_getnext_id()</i>	<i>setutxent()</i>
2470	<i>fgetpos()</i>	<i>getnetent()</i>	<i>posix_trace_eventtypelist_rewind()</i>	<i>strerror()</i>
2471	<i>fgets()</i>	<i>getprotobyname()</i>	<i>posix_trace_flush()</i>	<i>syslog()</i>
2472	<i>fgetwc()</i>	<i>getprotobynumber()</i>	<i>posix_trace_get_attr()</i>	<i>tmpfile()</i>
2473	<i>fgetws()</i>	<i>getprotoent()</i>	<i>posix_trace_get_filter()</i>	<i>tmpnam()</i>
2474	<i>fopen()</i>	<i>getpwent()</i>	<i>posix_trace_get_status()</i>	<i>ttyname()</i>
2475	<i>fprintf()</i>	<i>getpwnam()</i>	<i>posix_trace_getnext_event()</i>	<i>ttyname_r()</i>
2476	<i>fputc()</i>	<i>getpwnam_r()</i>	<i>posix_trace_open()</i>	<i>ungetc()</i>
2477	<i>fputs()</i>	<i>getpwuid()</i>	<i>posix_trace_rewind()</i>	<i>ungetwc()</i>
2478	<i>fputwc()</i>	<i>getpwuid_r()</i>	<i>posix_trace_set_filter()</i>	<i>unlink()</i>
2479	<i>fputws()</i>	<i>gets()</i>	<i>posix_trace_shutdown()</i>	<i>vfprintf()</i>
2480	<i>fread()</i>	<i>getservbyname()</i>	<i>posix_trace_timedgetnext_event()</i>	<i>vwprintf()</i>
2481	<i>freopen()</i>	<i>getservbyport()</i>	<i>posix_typed_mem_open()</i>	<i>vprintf()</i>
2482	<i>fscanf()</i>	<i>getservent()</i>	<i>printf()</i>	<i>vwprintf()</i>
2483	<i>fseek()</i>	<i>getutxent()</i>	<i>pthread_rwlock_rdlock()</i>	<i>wprintf()</i>
2484	<i>fseeko()</i>	<i>getutxid()</i>	<i>pthread_rwlock_timedrdlock()</i>	<i>wscanf()</i>
2485	<i>fsetpos()</i>	<i>getutxline()</i>	<i>pthread_rwlock_timedwrlock()</i>	

2486 An implementation shall not introduce cancelation points into any other functions specified in
 2487 this volume of IEEE Std. 1003.1-200x.

2488 _____
 2489 2. For any value of the *cmd* argument.

2490 The side effects of acting upon a cancelation request while suspended during a call of a function
2491 are the same as the side effects that may be seen in a single-threaded program when a call to
2492 a function is interrupted by a signal and the given function returns [EINTR]. Any such side effects
2493 occur before any cancelation cleanup handlers are called.

2494 Whenever a thread has cancelability enabled and a cancelation request has been made with that
2495 thread as the target and the thread calls *pthread_testcancel()*, then the cancelation request is acted
2496 upon before *pthread_testcancel()* returns. If a thread has cancelability enabled and the thread has
2497 a cancelation request pending and the thread is suspended at a cancelation point waiting for an
2498 event to occur, then the cancelation request shall be acted upon. However, if the thread is
2499 suspended at a cancelation point and the event that it is waiting for occurs before the cancelation
2500 request is acted upon, it is unspecified whether the cancelation request is acted upon or whether
2501 the request remains pending and the thread resumes normal execution.

2502 2.9.5.3 Thread Cancelation Cleanup Handlers

2503 Each thread maintains a list of cancelation cleanup handlers. The programmer uses the
2504 *pthread_cleanup_push()* and *pthread_cleanup_pop()* functions to place routines on and remove
2505 routines from this list.

2506 When a cancelation request is acted upon, the routines in the list are invoked one by one in LIFO
2507 sequence; that is, the last routine pushed onto the list (Last In) is the first to be invoked (First
2508 Out). The thread invokes the cancelation cleanup handler with cancelation disabled until the last
2509 cancelation cleanup handler returns. When the cancelation cleanup handler for a scope is
2510 invoked, the storage for that scope remains valid. If the last cancelation cleanup handler returns,
2511 thread execution is terminated and a status of PTHREAD_CANCELED is made available to any
2512 threads joining with the target. The symbolic constant PTHREAD_CANCELED expands to a
2513 constant expression of type (**void***) whose value matches no pointer to an object in memory nor
2514 the value NULL.

2515 The cancelation cleanup handlers are also invoked when the thread calls *pthread_exit()*.

2516 A side effect of acting upon a cancelation request while in a condition variable wait is that the
2517 mutex is re-acquired before calling the first cancelation cleanup handler. In addition, the thread
2518 is no longer considered to be waiting for the condition and the thread shall not have consumed
2519 any pending condition signals on the condition.

2520 A cancelation cleanup handler cannot exit via *longjmp()* or *siglongjmp()*.

2521 2.9.5.4 Async-Cancel Safety

2522 The *pthread_cancel()*, *pthread_setcancelstate()*, and *pthread_setcanceltype()* functions are defined to
2523 be async-cancel safe.

2524 No other functions in this volume of IEEE Std. 1003.1-200x are required to be async-cancel-safe.

2525 2.9.6 Thread Read-Write Locks

2526 Multiple readers, single writer (read-write) locks allow many threads to have simultaneous
2527 read-only access to data while allowing only one thread to have exclusive write access at any
2528 given time. They are typically used to protect data that is read-only more frequently than it is
2529 changed.

2530 One or more readers acquire read access to the resource by performing a read lock operation on
2531 the associated read-write lock. A writer acquires exclusive write access by performing a write
2532 lock operation. Basically, all readers exclude any writers and a writer excludes all readers and
2533 any other writers.

2534 A thread that has blocked on a read-write lock (for example, has not yet returned from a
2535 *pthread_rwlock_rdlock()* or *pthread_rwlock_wrlock()* call) shall not prevent any unblocked thread
2536 that is eligible to use the same processing resources from eventually making forward progress in
2537 its execution. Eligibility for processing resources shall be determined by the scheduling policy.

2538 Read-write locks can be used to synchronize threads in the current process and other processes if
2539 they are allocated in memory that is writable and shared among the cooperating processes and
2540 have been initialized for this behavior.

2541 **2.9.7 Thread Interactions with Regular File Operations**

2542 All of the functions *chmod()*, *close()*, *fchmod()*, *fcntl()*, *fstat()*, *ftruncate()*, *lseek()*, *open()*, *read()*,
2543 *readlink()*, *stat()*, *symlink()*, and *write()* shall be atomic with respect to each other in the effects
2544 specified in IEEE Std. 1003.1-200x when they operate on regular files. If two threads each call one
2545 of these functions, each call shall either see all of the specified effects of the other call, or none of
2546 them.

2547 **2.10 Sockets**

2548 A socket is an endpoint for communication using the facilities described in this section. A socket
2549 is created with a specific socket type, described in Section 2.10.6 (on page 563), and is associated
2550 with a specific protocol, detailed in Section 2.10.2. A socket is accessed via a file descriptor
2551 obtained when the socket is created.

2552 **2.10.1 Protocol Families**

2553 All network protocols are associated with a specific protocol family. A protocol family provides
2554 basic services to the protocol implementation to allow it to function within a specific network
2555 environment. These services may include packet fragmentation and reassembly, routing,
2556 addressing, and basic transport. A protocol family may support multiple methods of addressing.
2557 Each method represents an address family. A protocol family is normally comprised of a
2558 number of protocols, one per socket type. Each protocol is characterized by an abstract socket
2559 type. It is not required that a protocol family support all socket types. A protocol family may
2560 contain multiple protocols supporting the same socket abstraction.

2561 Section 2.10.17 (on page 569), Section 2.10.18 (on page 569), and Section 2.10.19 (on page 570),
2562 respectively, describe the use of sockets for local UNIX connections, for Internet protocols based
2563 on IPv4, and for Internet protocols based on IPv6.

2564 **2.10.2 Protocols**

2565 A protocol supports one of the socket abstractions detailed in Section 2.10.6 (on page 563).
2566 Selecting a protocol involves specifying the protocol family, socket type, and protocol number to
2567 the *socket()* function. Protocols normally accept only one type of address format, usually
2568 determined by the addressing structure inherent in the design of the protocol family/network
2569 architecture. Certain semantics of the basic socket abstractions are protocol-specific. All
2570 protocols are expected to support the basic model for their particular socket type, but may, in
2571 addition, provide non-standard facilities or extensions to a mechanism.

2572 **2.10.3 Addressing**

2573 Associated with each protocol family is at least one address family. An address family defines
2574 the format of a socket address. All network addresses are described using a general structure,
2575 called a **sockaddr**, as defined in the Base Definitions volume of IEEE Std. 1003.1-200x,
2576 <**sys/socket.h**>. However, each address family imposes finer and more specific structure,
2577 generally defining a structure with fields specific to the address family. The field *sa_family* in the
2578 **sockaddr** structure contains the address family identifier, specifying the format of the *sa_data*
2579 area. The size of the *sa_data* area is unspecified.

2580 **2.10.4 Routing**

2581 Sockets provides packet routing facilities. A routing information database is maintained, which
2582 is used in selecting the appropriate network interface when transmitting packets.

2583 2.10.5 Interfaces

2584 Each network interface in a system corresponds to a path through which messages can be sent
2585 and received. A network interface usually has a hardware device associated with it, though
2586 certain interfaces such as the loopback interface, do not.

2587 2.10.6 Socket Types

2588 A socket is created with a specific type, which defines the communication semantics and which
2589 allows the selection of an appropriate communication protocol. Three types are defined:
2590 SOCK_STREAM, SOCK_SEQPACKET, and SOCK_DGRAM. Implementations may specify
2591 additional socket types.

2592 The SOCK_STREAM socket type provides reliable, sequenced, full-duplex octet streams
2593 between the socket and a peer to which the socket is connected. A socket of type
2594 SOCK_STREAM must be in a connected state before any data may be sent or received. Record
2595 boundaries are not maintained; data sent on a stream socket using output operations of one size
2596 may be received using input operations of smaller or larger sizes without loss of data. Data may
2597 be buffered; successful return from an output function does not imply that the data has been
2598 delivered to the peer or even transmitted from the local system. If data cannot be successfully
2599 transmitted within a given time then the connection is considered broken, and subsequent
2600 operations shall fail. A SIGPIPE signal is raised if a thread sends on a broken stream (one that is
2601 no longer connected). Support for an out-of-band data transmission facility is protocol-specific.

2602 The SOCK_SEQPACKET socket type is similar to the SOCK_STREAM type, and is also
2603 connection-oriented. The only difference between these types is that record boundaries are
2604 maintained using the SOCK_SEQPACKET type. A record can be sent using one or more output
2605 operations and received using one or more input operations, but a single operation never
2606 transfers parts of more than one record. Record boundaries are visible to the receiver via the
2607 MSG_EOR flag in the received message flags returned by the *recvmsg()* function. It is protocol-
2608 specific whether a maximum record size is imposed.

2609 The SOCK_DGRAM socket type supports connectionless data transfer which is not necessarily
2610 acknowledged or reliable. Datagrams may be sent to a peer named in each output operation, and
2611 incoming datagrams may be received from multiple sources. The source address of each
2612 datagram is available when receiving the datagram. An application may also pre-specify a peer
2613 address, in which case calls to output functions shall send to the pre-specified peer. If a peer has
2614 been specified, only datagrams from that peer shall be received. A datagram must be sent in a
2615 single output operation, and must be received in a single input operation. The maximum size of
2616 a datagram is protocol-specific; with some protocols, the limit is implementation-defined.
2617 Output datagrams may be buffered within the system; thus, a successful return from an output
2618 function does not guarantee that a datagram is actually sent or received. However,
2619 implementations should attempt to detect any errors possible before the return of an output
2620 function, reporting any error by an unsuccessful return value.

2621 2.10.7 Socket I/O Mode

2622 The I/O mode of a socket is described by the O_NONBLOCK file status flag which pertains to
2623 the open file description for the socket. This flag is initially off when a socket is created, but may
2624 be set and cleared by the use of the F_SETFL command of the *fcntl()* function.

2625 When the O_NONBLOCK flag is set, functions that would normally block until they are
2626 complete either return immediately with an error, or they complete asynchronously to the
2627 execution of the calling process. Data transfer operations (the *read()*, *write()*, *send()*, and *recv()*
2628 functions) complete immediately, transfer only as much as is available, and then return without
2629 blocking, or return an error indicating that no transfer could be made without blocking. The

2630 *connect()* function initiates a connection and returns without blocking when `O_NONBLOCK` is
2631 set; it returns the error `[EINPROGRESS]` to indicate that the connection was initiated
2632 successfully, but that it has not yet completed.

2633 **2.10.8 Socket Owner**

2634 The owner of a socket is unset when a socket is created. The owner may be set to a process ID or
2635 process group ID using the `F_SETOWN` command of the *fcntl()* function.

2636 **2.10.9 Socket Queue Limits**

2637 The transmit and receive queue sizes for a socket are set when the socket is created. The default
2638 sizes used are both protocol-specific and implementation-defined. The sizes may be changed
2639 using the *setsockopt()* function.

2640 **2.10.10 Pending Error**

2641 Errors may occur asynchronously, and be reported to the socket in response to input from the
2642 network protocol. The socket stores the pending error to be reported to a user of the socket at the
2643 next opportunity. The error is returned in response to a subsequent *send()*, *recv()*, or *getsockopt()*
2644 operation on the socket, and the pending error is then cleared.

2645 **2.10.11 Socket Receive Queue**

2646 A socket has a receive queue that buffers data when they are received by the system until they
2647 are removed by a receive call. Depending on the type of the socket and the communication
2648 provider, the receive queue may also contain ancillary data such as the addressing and other
2649 protocol data associated with the normal data in the queue, and may contain out-of-band or
2650 expedited data. The limit on the queue size includes any normal, out-of-band data, datagram
2651 source addresses, and ancillary data in the queue. The description in this section applies to all
2652 sockets, even though some elements cannot be present in some instances.

2653 The contents of a receive buffer are logically structured as a series of data segments with
2654 associated ancillary data and other information. A data segment may contain normal data or
2655 out-of-band data, but never both. A data segment may complete a record if the protocol
2656 supports records (always true for types `SOCK_SEQPACKET` and `SOCK_DGRAM`). A record
2657 may be stored as more than one segment; the complete record might never be present in the
2658 receive buffer at one time, as a portion might already have been returned to the application, and
2659 another portion might not yet have been received from the communications provider. A data
2660 segment may contain ancillary protocol data, which is logically associated with the segment.
2661 Ancillary data is received as if it were queued along with the first normal data octet in the
2662 segment (if any). A segment may contain ancillary data only, with no normal or out-of-band
2663 data. For the purposes of this section, a datagram is considered to be a data segment that
2664 terminates a record, and that includes a source address as a special type of ancillary data. Data
2665 segments are placed into the queue as data is delivered to the socket by the protocol. Normal
2666 data segments are placed at the end of the queue as they are delivered. If a new segment
2667 contains the same type of data as the preceding segment and includes no ancillary data, and if
2668 the preceding segment does not terminate a record, the segments are logically merged into a
2669 single segment.

2670 The receive queue is logically terminated if an end-of-file indication has been received or a
2671 connection has been terminated. A segment shall be considered to be terminated if another
2672 segment follows it in the queue, if the segment completes a record, or if an end-of-file or other
2673 connection termination has been reported. The last segment in the receive queue shall also be
2674 considered to be terminated while the socket has a pending error to be reported.

2675 A receive operation shall never return data or ancillary data from more than one segment.

2676 **2.10.12 Socket Out-of-Band Data State**

2677 The handling of received out-of-band data is protocol-specific. Out-of-band data may be placed
 2678 in the socket receive queue, either at the end of the queue or before all normal data in the queue.
 2679 In this case, out-of-band data is returned to an application program by a normal receive call.
 2680 Out-of-band data may also be queued separately rather than being placed in the socket receive
 2681 queue, in which case it shall be returned only in response to a receive call that requests out-of-
 2682 band data. It is protocol-specific whether an out-of-band data mark is placed in the receive
 2683 queue to demarcate data preceding the out-of-band data and following the out-of-band data. An
 2684 out-of-band data mark is logically an empty data segment that cannot be merged with other
 2685 segments in the queue. An out-of-band data mark is never returned in response to an input
 2686 operation. The *socketmark()* function can be used to test whether an out-of-band data mark is the
 2687 first element in the queue. If an out-of-band data mark is the first element in the queue when an
 2688 input function is called without the MSG_PEEK option, the mark is removed from the queue and
 2689 the following data (if any) are processed as if the mark had not been present.

2690 **2.10.13 Connection Indication Queue**

2691 Sockets that are used to accept incoming connections maintain a queue of outstanding
 2692 connection indications. This queue is a list of connections that are awaiting acceptance by the
 2693 application. See *listen()*.

2694 **2.10.14 Signals**

2695 One category of event at the socket interface is the generation of signals. These signals report
 2696 protocol events or process errors relating to the state of the socket. The generation or delivery of
 2697 a signal does not change the state of the socket, although the generation of the signal may have
 2698 been caused by a state change.

2699 The SIGPIPE signal shall be sent to a thread that attempts to send data on a socket that is no
 2700 longer able to send. In addition, the send operation fails with the error [EPIPE].

2701 If a socket has an owner, the SIGURG signal is sent to the owner of the socket when it is notified
 2702 of expedited or out-of-band data. The socket state at this time is protocol-dependent, and the
 2703 status of the socket is specified in Section 2.10.17 (on page 569), Section 2.10.18 (on page 569),
 2704 and Section 2.10.19 (on page 570). Depending on the protocol, the expedited data may or may
 2705 not have arrived at the time of signal generation.

2706 **2.10.15 Asynchronous Errors**

2707 If any of the following conditions occur asynchronously for a socket, the corresponding value
 2708 listed below shall become the pending error for the socket:

2709 [ECONNABORTED]

2710 The connection was aborted locally.

2711 [ECONNREFUSED]

2712 For a connection-mode socket attempting a non-blocking connection, the attempt to connect
 2713 was forcefully rejected. For a connectionless-mode socket, an attempt to deliver a datagram
 2714 was forcefully rejected.

2715 [ECONNRESET]

2716 The peer has aborted the connection.

- 2717 [EHOSTDOWN]
 2718 The destination host has been determined to be down or disconnected.
- 2719 [EHOSTUNREACH]
 2720 The destination host is not reachable.
- 2721 [EMSGSIZE]
 2722 For a connectionless-mode socket, the size of a previously sent datagram prevented
 2723 delivery.
- 2724 [ENETDOWN]
 2725 The local network connection is not operational.
- 2726 [ENETRESET]
 2727 The connection was aborted by the network.
- 2728 [ENETUNREACH]
 2729 The destination network is not reachable.

2730 2.10.16 Use of Options

2731 There are a number of socket options which either specialize the behavior of a socket or provide
 2732 useful information. These options may be set at different protocol levels and are always present
 2733 at the uppermost “socket” level.

2734 Socket options are manipulated by two functions, *getsockopt()* and *setsockopt()*. These functions
 2735 allow an application program to customize the behavior and characteristics of a socket to
 2736 provide the desired effect.

2737 All of the options have default values. The type and meaning of these values is defined by the
 2738 protocol level to which they apply. Instead of using the default values, an application program
 2739 may choose to customize one or more of the options. However, in the bulk of cases, the default
 2740 values are sufficient for the application.

2741 Some of the options are used to enable or disable certain behavior within the protocol modules
 2742 (for example, turn on debugging) while others may be used to set protocol-specific information
 2743 (for example, IP time-to-live on all the application’s outgoing packets). As each of the options is
 2744 introduced, its effect on the underlying protocol modules is described.

2745 Table 2-4 shows the value for the socket level.

2746 **Table 2-4** Value of Level for Socket Options

Name	Description
SOL_SOCKET	Options are intended for the sockets level.

2749 Table 2-5 (on page 567) lists those options present at the socket level; that is, when the *level*
 2750 parameter of the *getsockopt()* or *setsockopt()* function is SOL_SOCKET, the types of the option
 2751 value parameters associated with each option, and a brief synopsis of the meaning of the option
 2752 value parameter. Unless otherwise noted, each may be examined with *getsockopt()* and set with
 2753 *setsockopt()* on all types of socket.

2754

Table 2-5 Socket-Level Options

2755

2756

2757

2758

2759

2760

2761

2762

2763

2764

2765

2766

2767

2768

2769

2770

2771

2772

2773

2774

2775

2776

2777

2778

2779

2780

2781

2782

Option	Parameter Type	Parameter Meaning
SO_BROADCAST	(void *) int	Non-zero requests permission to transmit broadcast datagrams (SOCK_DGRAM sockets only).
SO_DEBUG	int	Non-zero requests debugging in underlying protocol modules.
SO_DONTROUTE	int	Non-zero requests bypass of normal routing; route based on destination address only.
SO_ERROR	int	Requests and clears pending error information on the socket (<i>getsockopt()</i> only).
SO_KEEPALIVE	int	Non-zero requests periodic transmission of keepalive messages (protocol-specific).
SO_LINGER	struct linger	Specify actions to be taken for queued, unsent data on <i>close()</i> : linger on/off and linger time in seconds.
SO_OOBINLINE	int	Non-zero requests that out-of-band data be placed into normal data input queue as received.
SO_RCVBUF	int	Size of receive buffer (in bytes).
SO_RCVLOWAT	int	Minimum amount of data to return to application for input operations (in bytes).
SO_RCVTIMEO	struct timeval	Timeout value for a socket receive operation.
SO_REUSEADDR	int	Non-zero requests reuse of local addresses in <i>bind()</i> (protocol-specific).
SO_SNDBUF	int	Size of send buffer (in bytes).
SO_SNDLOWAT	int	Minimum amount of data to send for output operations (in bytes).
SO_SNDTIMEO	struct timeval	Timeout value for a socket send operation.
SO_TYPE	int	Identify socket type (<i>getsockopt()</i> only).

2783

2784

2785

The SO_BROADCAST option requests permission to send broadcast datagrams on the socket. Support for SO_BROADCAST is protocol-specific. The default for SO_BROADCAST is that the ability to send broadcast datagrams on a socket is disabled.

2786

2787

2788

2789

SO_DEBUG enables debugging in the underlying protocol modules. This can be useful for tracing the behavior of the underlying protocol modules during normal system operation. The semantics of the debug reports are implementation-defined. The default value for SO_DEBUG is for debugging to be turned off.

2790

2791

2792

2793

2794

SO_DONTROUTE requests that outgoing messages bypass the standard routing facilities. The destination must be on a directly-connected network, and messages are directed to the appropriate network interface according to the destination address. It is protocol-specific whether this option has any effect and how the outgoing network interface is chosen. Support for this option with each protocol is implementation-defined.

2795

2796

2797

2798

2799

SO_ERROR is used only on *getsockopt()*. When this option is specified, *getsockopt()* returns any pending error on the socket and clears the error status. It returns a value of 0 if there is no pending error. SO_ERROR may be used to check for asynchronous errors on connected connectionless-mode sockets or for other types of asynchronous errors. SO_ERROR has no default value.

2800 SO_KEEPALIVE enables the periodic transmission of messages on a connected socket. The
2801 behavior of this option is protocol-specific. The default value for SO_KEEPALIVE is zero,
2802 specifying that this capability is turned off.

2803 The SO_LINGER option controls the action of the interface when unsent messages are queued
2804 on a socket and a *close()* is performed. The details of this option are protocol-specific. The
2805 default value for SO_LINGER is zero, or off, for the *L_onoff* element of the option value and zero
2806 seconds for the linger time specified by the *L_linger* element.

2807 SO_OOBINLINE is valid only on protocols that support out-of-band data. The SO_OOBINLINE
2808 option requests that out-of-band data be placed in the normal data input queue as received; it is
2809 then accessible using the *read()* or *recv()* functions without the MSG_OOB flag set. The default
2810 for SO_OOBINLINE is off; that is, for out-of-band data not to be placed in the normal data input
2811 queue.

2812 SO_RCVBUF requests that the buffer space allocated for receive operations on this socket be set
2813 to the value, in bytes, of the option value. Applications may wish to increase buffer size for high
2814 volume connections, or may decrease buffer size to limit the possible backlog of incoming data.
2815 The default value for the SO_RCVBUF option value is implementation-defined, and may vary by
2816 protocol.

2817 The maximum value for the option for a socket may be obtained by the use of the *fpathconf()*
2818 function, using the value *_PC_SOCKET_MAXBUF*.

2819 SO_RCVLOWAT sets the minimum number of bytes to process for socket input operations. In
2820 general, receive calls block until any (non-zero) amount of data is received, then return the
2821 smaller of the amount available or the amount requested. The default value for SO_RCVLOWAT
2822 is 1, and does not affect the general case. If SO_RCVLOWAT is set to a larger value, blocking
2823 receive calls normally wait until they have received the smaller of the low water mark value or
2824 the requested amount. Receive calls may still return less than the low water mark if an error
2825 occurs, a signal is caught, or the type of data next in the receive queue is different than that
2826 returned (for example, out-of-band data). As mentioned previously, the default value for
2827 SO_RCVLOWAT is 1 byte. It is implementation-defined whether the SO_RCVLOWAT option
2828 can be set.

2829 SO_RCVTIMEO is an option to set a timeout value for input operations. It accepts a **timeval**
2830 structure with the number of seconds and microseconds specifying the limit on how long to wait
2831 for an input operation to complete. If a receive operation has blocked for this much time without
2832 receiving additional data, it returns with a partial count or *errno* set to [EWOULDBLOCK] if no
2833 data were received. The default for this option is the value zero, which indicates that a receive
2834 operation will not timeout. It is implementation-defined whether the SO_RCVTIMEO option can
2835 be set.

2836 SO_REUSEADDR indicates that the rules used in validating addresses supplied in a *bind()*
2837 should allow reuse of local addresses. Operation of this option is protocol-specific. The default
2838 value for SO_REUSEADDR is off; that is, reuse of local addresses is not permitted.

2839 SO_SNDBUF requests that the buffer space allocated for send operations on this socket be set to
2840 the value, in bytes, of the option value. The default value for the SO_SNDBUF option value is
2841 implementation-defined, and may vary by protocol. The maximum value for the option for a
2842 socket may be obtained by the use of the *fpathconf()* function, using the value
2843 *_PC_SOCKET_MAXBUF*.

2844 SO_SNDLOWAT sets the minimum number of bytes to process for socket output operations.
2845 Most output operations process all of the data supplied by the call, delivering data to the
2846 protocol for transmission and blocking as necessary for flow control. Non-blocking output
2847 operations process as much data as permitted subject to flow control without blocking, but

2848 process no data if flow control does not allow the smaller of the send low water mark value or
 2849 the entire request to be processed. A *select()* operation testing the ability to write to a socket
 2850 returns true only if the send low water mark could be processed. The default value for
 2851 SO_SNDLOWAT is implementation-defined and protocol-specific. It is implementation-defined
 2852 whether the SO_SNDLOWAT option can be set.

2853 SO_SNDTIMEO is an option to set a timeout value for the amount of time that an output
 2854 function shall block because flow control prevents data from being sent. As noted in Table 2-5
 2855 (on page 567), the option value is a **timeval** structure with the number of seconds and
 2856 microseconds specifying the limit on how long to wait for an output operation to complete. If a
 2857 send operation has blocked for this much time, it returns with a partial count or *errno* set to
 2858 [EWOULDBLOCK] if no data were sent. The default for this option is the value zero, which
 2859 indicates that a send operation will not timeout. It is implementation-defined whether the
 2860 SO_SNDTIMEO option can be set.

2861 SO_TYPE is used only on *getsockopt()*. When this option is specified, *getsockopt()* returns the
 2862 type of the socket (for example, SOCK_STREAM). This option is useful to servers that inherit
 2863 sockets on start-up. SO_TYPE has no default value.

2864 **2.10.17 Use of Sockets for Local UNIX Connections**

2865 Support for UNIX domain sockets is mandatory.

2866 UNIX domain sockets provide process-to-process communication in a single system.

2867 *2.10.17.1 Headers*

2868 Symbolic constant AF_UNIX is defined in the `<sys/socket.h>` header to identify the UNIX
 2869 domain address family. The `<sys/un.h>` header contains other definitions used in connection
 2870 with UNIX domain sockets. See the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter
 2871 13, Headers.

2872 The **sockaddr_storage** structure defined in `<sys/socket.h>` is large enough to accommodate a
 2873 **sockaddr_un** structure (see the `<sys/un.h>` header defined in the Base Definitions volume of
 2874 IEEE Std. 1003.1-200x, Chapter 13, Headers) and is aligned at an appropriate boundary so that
 2875 pointers to it can be cast as pointers to **sockaddr_un** structures and used to access the fields of
 2876 those structures without alignment problems. When a **sockaddr_storage** structure is cast as a
 2877 **sockaddr_un** structure, the *ss_family* field maps onto the *sun_family* field.

2878 **2.10.18 Use of Sockets over Internet Protocols Based on IPv4**

2879 Support for sockets over Internet protocols based on IPv4 is mandatory.

2880 *2.10.18.1 Headers*

2881 Symbolic constant AF_INET is defined in the `<sys/socket.h>` header to identify the IPv4 Internet
 2882 address family. The `<netinet/in.h>` header contains other definitions used in connection with
 2883 IPv4 Internet sockets. See the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 13,
 2884 Headers.

2885 The **sockaddr_storage** structure defined in `<sys/socket.h>` is large enough to accommodate a
 2886 **sockaddr_in** structure (see the `<netinet/in.h>` header defined in the Base Definitions volume of
 2887 IEEE Std. 1003.1-200x, Chapter 13, Headers) and is aligned at an appropriate boundary so that
 2888 pointers to it can be cast as pointers to **sockaddr_in** structures and used to access the fields of
 2889 those structures without alignment problems. When a **sockaddr_storage** structure is cast as a
 2890 **sockaddr_in** structure, the *ss_family* field maps onto the *sin_family* field.

2891 2.10.19 Use of Sockets over Internet Protocols Based on IPv6

2892 IP6 This section describes extensions to support sockets over Internet protocols based on IPv6. This
2893 functionality is dependent on support of the IPV6 option (and the rest of this section is not
2894 further shaded for this option).

2895 To enable smooth transition from IPv4 to IPv6, the features defined in this section may, in certain
2896 circumstances, also be used in connection with IPv4; see Section 2.10.19.2 (on page 571).

2897 2.10.19.1 Addressing

2898 IPv6 overcomes the addressing limitations of previous versions by using 128-bit addresses
2899 instead of 32-bit addresses. The IPv6 address architecture is described in RFC 2373.

2900 There are three kinds of IPv6 address:

2901 Unicast

2902 Identifies a single interface.

2903 A unicast address can be global, link-local (designed for use on a single link), or site-local
2904 (designed for systems not connected to the Internet). Link-local and site-local addresses
2905 need not be globally unique.

2906 Anycast

2907 Identifies a set of interfaces such that a packet sent to the address can be delivered to any
2908 member of the set.

2909 An anycast address is similar to a unicast address; the nodes to which an anycast address is
2910 assigned must be explicitly configured to know that it is an anycast address.

2911 Multicast

2912 Identifies a set of interfaces such that a packet sent to the address should be delivered to
2913 every member of the set.

2914 An application can send multicast datagrams by simply specifying an IPv6 multicast
2915 address in the *address* argument of *sendto()*. To receive multicast datagrams, an application
2916 must join the multicast group (using *setsockopt()* with IPV6_JOIN_GROUP) and must bind
2917 to the socket the UDP port on which datagrams will be received. Some applications should
2918 also bind the multicast group address to the socket, to prevent other datagrams destined to
2919 that port from being delivered to the socket.

2920 A multicast address can be global, node-local, link-local, site-local, or organization-local.

2921 The following special IPv6 addresses are defined:

2922 Unspecified

2923 An address that is not assigned to any interface and is used to indicate the absence of an
2924 address.

2925 Loopback

2926 A unicast address that is not assigned to any interface and can be used by a node to send
2927 packets to itself.

2928 Two sets of IPv6 addresses are defined to correspond to IPv4 addresses:

2929 IPv4-compatible addresses

2930 These are assigned to nodes that support IPv6 and can be used when traffic is “tunneled”
2931 through IPv4.

2932 IPv4-mapped addresses

2933 These are used to represent IPv4 addresses in IPv6 address format; see Section 2.10.19.2 (on

2934 page 571).

2935 Note that the unspecified address and the loopback address must not be treated as IPv4-
2936 compatible addresses.

2937 2.10.19.2 Compatibility with IPv4

2938 The API provides the ability for IPv6 applications to interoperate with applications using IPv4,
2939 by using IPv4-mapped IPv6 addresses. These addresses can be generated automatically by the
2940 *getipnodebyname()* function when the specified host has only IPv4 addresses (as described in
2941 *endhostent()*).

2942 Applications may use AF_INET6 sockets to open TCP connections to IPv4 nodes, or send UDP
2943 packets to IPv4 nodes, by simply encoding the destination's IPv4 address as an IPv4-mapped
2944 IPv6 address, and passing that address, within a **sockaddr_in6** structure, in the *connect()*,
2945 *sendto()* or *sendmsg()* function. When applications use AF_INET6 sockets to accept TCP
2946 connections from IPv4 nodes, or receive UDP packets from IPv4 nodes, the system returns the
2947 peer's address to the application in the *accept()*, *recvfrom()*, *recvmsg()*, or *getpeername()*
2948 function using a **sockaddr_in6** structure encoded this way. If a node has an IPv4 address, then the
2949 implementation may allow applications to communicate using that address via an AF_INET6
2950 socket. In such a case, the address will be represented at the API by the corresponding IPv4-
2951 mapped IPv6 address. Also, the implementation may allow an AF_INET6 socket bound to
2952 **in6addr_any** to receive inbound connections and packets destined to one of the node's IPv4
2953 addresses.

2954 An application may use AF_INET6 sockets to bind to a node's IPv4 address by specifying the
2955 address as an IPv4-mapped IPv6 address in a **sockaddr_in6** structure in the *bind()* function. For
2956 an AF_INET6 socket bound to a node's IPv4 address, the system returns the address in the
2957 *getsockname()* function as an IPv4-mapped IPv6 address in a **sockaddr_in6** structure.

2958 2.10.19.3 Interface Identification

2959 Each local interface is assigned a unique positive integer as a numeric index. Indexes start at 1;
2960 zero is not used. There may be gaps so that there is no current interface for a particular positive
2961 index. Each interface also has a unique implementation-defined name.

2962 2.10.19.4 Options

2963 The following options apply at the IPPROTO_IPV6 level:

2964 IPV6_JOIN_GROUP

2965 When set via *setsockopt()*, it joins the application to a multicast group on an interface
2966 (identified by its index) and addressed by a given multicast address, enabling packets sent
2967 to that address to be read via the socket. If the interface index is specified as zero, the
2968 system selects the interface (for example, by looking up the address in a routing table and
2969 using the resulting interface).

2970 An attempt to read this option using *getsockopt()* results in an [EOPNOTSUPP] error.

2971 The value of this option is an **ipv6_mreq** structure.

2972 IPV6_LEAVE_GROUP

2973 When set via *setsockopt()*, it removes the application from the multicast group on an
2974 interface (identified by its index) and addressed by a given multicast address.

2975 An attempt to read this option using *getsockopt()* results in an [EOPNOTSUPP] error.

- 2976 The value of this option is an **ipv6_mreq** structure.
- 2977 **IPV6_MULTICAST_HOPS**
- 2978 The value of this option is the hop limit for outgoing multicast IPv6 packets sent via the
2979 socket. Its possible values are the same as those of **IPV6_UNICAST_HOPS**. If the
2980 **IPV6_MULTICAST_HOPS** option is not set, a value of 1 is assumed. This option can be set
2981 via *setsockopt()* and read via *getsockopt()*.
- 2982 **IPV6_MULTICAST_IF**
- 2983 The index of the interface to be used for outgoing multicast packets. It can be set via
2984 *setsockopt()* and read via *getsockopt()*.
- 2985 **IPV6_MULTICAST_LOOP**
- 2986 This option controls whether outgoing multicast packets should be delivered back to the
2987 local application when the sending interface is itself a member of the destination multicast
2988 group. If it is set to 1 they are delivered. If it is set to 0 they are not. Other values result in an
2989 [EINVAL] error. This option can be set via *setsockopt()* and read via *getsockopt()*.
- 2990 **IPV6_UNICAST_HOPS**
- 2991 The value of this option is the hop limit for outgoing unicast IPv6 packets sent via the
2992 socket. If the option is not set, or is set to -1, the system selects a default value. Attempts to
2993 set a value less than -1 or greater than 255 result in an [EINVAL] error. This option can be
2994 set via *setsockopt()* and read via *getsockopt()*.
- 2995 An [EOPNOTSUPP] error results if **IPV6_JOIN_GROUP** or **IPV6_LEAVE_GROUP** is used with
2996 *getsockopt()*.
- 2997 2.10.19.5 Headers
- 2998 Symbolic constant **AF_INET6** is defined in the `<sys/socket.h>` header to identify the IPv6
2999 Internet address family. See the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 13,
3000 Headers.
- 3001 The **sockaddr_storage** structure defined in `<sys/socket.h>` is large enough to accommodate a
3002 **sockaddr_in6** structure (see the `<netinet/in.h>` header defined in the Base Definitions volume of
3003 IEEE Std. 1003.1-200x, Chapter 13, Headers) and is aligned at an appropriate boundary so that
3004 pointers to it can be cast as pointers to **sockaddr_in6** structures and used to access the fields of
3005 those structures without alignment problems. When a **sockaddr_storage** structure is cast as a
3006 **sockaddr_in6** structure, the *ss_family* field maps onto the *sin6_family* field.
- 3007 The `<netinet/in.h>`, `<arpa/inet.h>`, and `<netdb.h>` headers contain other definitions used in
3008 connection with IPv6 Internet sockets; see the Base Definitions volume of IEEE Std. 1003.1-200x,
3009 Chapter 13, Headers.

3010 2.11 Tracing

3011 TRC This section describes extensions to support tracing of user applications. This functionality is
3012 dependent on support of the Trace option (and the rest of this section is not further shaded for
3013 this option).

3014 The tracing facilities defined in IEEE Std. 1003.1-200x allow a process to select a set of trace event
3015 types, to activate a trace stream of the selected trace events as they occur in the flow of
3016 execution, and to retrieve the recorded trace events.

3017 The tracing operation relies on three logically different components: the traced process, the
3018 controller process, and the analyzer process. During the execution of the traced process, when a
3019 trace point is reached, a trace event is recorded into the trace streams created for that process in
3020 which the associated trace event type identifier is not being filtered out. The controller process
3021 controls the operation of recording the trace events into the trace stream. It shall be able to:

- 3022 • Initialize the attributes of a trace stream
- 3023 • Create the trace stream (for a specified traced process) using those attributes
- 3024 • Start and stop tracing for the trace stream
- 3025 • Filter the type of trace events to be recorded, if the Trace Event Filter option is supported
- 3026 • Shut a trace stream down

3027 These operations can be done for an active trace stream. The analyzer process retrieves the
3028 traced events either at runtime, when the trace stream has not yet been shut down, but is still
3029 recording trace events; or after opening a trace log that had been previously recorded and shut
3030 down. These three logically different operations can be performed by the same process, or can be
3031 distributed into different processes.

3032 A trace stream identifier can be created by a call to *posix_trace_create()*,
3033 *posix_trace_create_withlog()*, or *posix_trace_open()*. The *posix_trace_create()* and
3034 *posix_trace_create_withlog()* functions should be used by a controller process. The
3035 *posix_trace_open()* should be used by an analyzer process.

3036 The tracing functions can serve different purposes. One purpose is debugging the possibly pre-
3037 instrumented code, while another is post-mortem fault analysis. These two potential uses differ
3038 in that the first requires pre-filtering capabilities to avoid overwhelming the trace stream and
3039 permits focusing on expected information; while the second needs comprehensive trace
3040 capabilities in order to be able to record all types of information.

3041 The events to be traced belong to two classes:

- 3042 1. User trace events (generated by the application instrumentation)
- 3043 2. System trace events (generated by the operating system)

3044 The trace interface defines several system trace event types associated with control of and
3045 operation of the trace stream. This small set of system trace events includes the minimum
3046 required to interpret correctly the trace event information present in the stream. Other desirable
3047 system trace events for some particular application profile may be implemented and are
3048 encouraged; for example, process and thread scheduling, signal occurrence, and so on.

3049 Each traced process shall have a mapping of the trace event names to trace event type identifiers
3050 that have been defined for that process. Each active trace stream shall have a mapping that
3051 incorporates all the trace event type identifiers predefined by the trace system plus all the
3052 mappings of trace event names to trace event type identifiers of the processes that are being
3053 traced into that trace stream. These mappings are defined from the instrumented application by

3054 calling the *posix_trace_eventid_open()* function and from the controller process by calling the
 3055 *posix_trace_trid_eventid_open()* function. For a pre-recorded trace stream, the list of trace event
 3056 types is obtained from the pre-recorded trace log.

3057 The *st_ctime* and *st_mtime* fields of a file associated with an active trace stream shall be marked
 3058 for update every time any of the tracing operations modifies that file.

3059 The *st_atime* field of a file associated with a trace stream shall be marked for update every time
 3060 any of the tracing operations causes data to be read from that file.

3061 Results are undefined if the application performs any operation on a file descriptor associated
 3062 with an active or pre-recorded trace stream until *posix_trace_shutdown()* or *posix_trace_close()* is
 3063 called for that trace stream.

3064 The main purpose of this option is to define a complete set of functions and concepts that allow
 3065 a portable application to be traced from birth to death, whatever its realtime constraints and
 3066 properties.

3067 2.11.1 Tracing Data Definitions

3068 2.11.1.1 Structures

3069 The **<trace.h>** header shall define the *posix_trace_status_info* and *posix_trace_event_info* structures
 3070 described below. Implementations may add extensions to these structures.

3071 **posix_trace_status_info** Structure

3072 To facilitate control of a trace stream, information about the current state of an active trace
 3073 stream can be obtained dynamically. This structure is returned by a call to the
 3074 *posix_trace_get_status()* function.

3075 The **posix_trace_status_info** structure defined in **<trace.h>** shall contain at least the following
 3076 members:

3077

3078

3079

3080

3081

3082

Member Type	Member Name	Description
int	<i>posix_stream_status</i>	The operating mode of the trace stream.
int	<i>posix_stream_full_status</i>	The full status of the trace stream.
int	<i>posix_stream_overrun_status</i>	Indicates whether trace events were lost in the trace stream.

3083 If the Trace Log option is supported in addition to the Trace option, the **posix_trace_status_info**
 3084 structure defined in **<trace.h>** shall contain at least the following additional members:

3085

3086

3087

3088

3089

3090

3091

3092

Member Type	Member Name	Description
int	<i>posix_stream_flush_status</i>	Indicates whether a flush is in progress.
int	<i>posix_stream_flush_error</i>	Indicates whether any error occurred during the last flush operation.
int	<i>posix_log_overrun_status</i>	Indicates whether trace events were lost in the trace log.
int	<i>posix_log_full_status</i>	The full status of the trace log.

3093 The *posix_stream_status* member indicates the operating mode of the trace stream and shall have
 3094 one of the following values defined by manifest constants in the **<trace.h>** header:

- 3095 POSIX_TRACE_RUNNING
3096 Tracing is in progress; that is, the trace stream is accepting trace events.
- 3097 POSIX_TRACE_SUSPENDED
3098 The trace stream is not accepting trace events. The tracing operation has not yet started or
3099 has stopped, either following a *posix_trace_stop()* function call or because the trace resources
3100 are exhausted.
- 3101 The *posix_stream_full_status* member indicates the full status of the trace stream, and it shall have
3102 one of the following values defined by manifest constants in the `<trace.h>` header:
- 3103 POSIX_TRACE_FULL
3104 The space in the trace stream for trace events is exhausted.
- 3105 POSIX_TRACE_NOT_FULL
3106 There is still space available in the trace stream.
- 3107 The combination of the *posix_stream_status* and *posix_stream_full_status* members also indicates
3108 the actual status of the stream. The status shall be interpreted as follows:
- 3109 POSIX_TRACE_RUNNING and POSIX_TRACE_NOT_FULL
3110 This status combination indicates that tracing is in progress, and there is space available for
3111 recording more trace events.
- 3112 POSIX_TRACE_RUNNING and POSIX_TRACE_FULL
3113 This status combination indicates that tracing is in progress and that the trace stream is full
3114 of trace events. This status combination cannot occur unless the *stream-full-policy* is set to
3115 POSIX_TRACE_LOOP. The trace stream contains trace events recorded during a moving
3116 time window of prior trace events, and some older trace events may have been overwritten
3117 and thus lost.
- 3118 POSIX_TRACE_SUSPENDED and POSIX_TRACE_NOT_FULL
3119 This status combination indicates that tracing has not yet been started, has been stopped by
3120 the *posix_trace_stop()* function, or has been cleared by the *posix_trace_clear()* function.
- 3121 POSIX_TRACE_SUSPENDED and POSIX_TRACE_FULL
3122 This status combination indicates that tracing has been stopped by the implementation
3123 because the *stream-full-policy* attribute was POSIX_TRACE_UNTIL_FULL and trace
3124 resources were exhausted, or that the trace stream was stopped by the function
3125 *posix_trace_stop()* at a time when trace resources were exhausted.
- 3126 The *posix_stream_overrun_status* member indicates whether trace events were lost in the trace
3127 stream, and shall have one of the following values defined by manifest constants in the
3128 `<trace.h>` header:
- 3129 POSIX_TRACE_OVERRUN
3130 At least one trace event was lost and thus was not recorded in the trace stream.
- 3131 POSIX_TRACE_NO_OVERRUN
3132 No trace events were lost.
- 3133 When the corresponding trace stream is created, the *posix_stream_overrun_status* member shall be
3134 set to POSIX_TRACE_NO_OVERRUN.
- 3135 Whenever an overrun occurs, *posix_stream_overrun_status* member shall be set to
3136 POSIX_TRACE_OVERRUN.
- 3137 An overrun occurs when:

- 3138 • The policy is `POSIX_TRACE_LOOP` and a recorded trace event is overwritten.
- 3139 • The policy is `POSIX_TRACE_UNTIL_FULL` and the trace stream is full when a trace event is
- 3140 generated.
- 3141 • If the Trace Log option is supported, the policy is `POSIX_TRACE_FLUSH` and at least one
- 3142 trace event is lost while flushing the trace stream to the trace log.
- 3143 The *posix_stream_overrun_status* member is reset to zero after its value is read.
- 3144 If the Trace Log option is supported in addition to the Trace option, the *posix_stream_flush_status*,
- 3145 *posix_stream_flush_error*, *posix_log_overrun_status*, and *posix_log_full_status* members are defined
- 3146 as follows; otherwise, they are undefined.
- 3147 The *posix_stream_flush_status* member indicates whether a flush operation is being performed
- 3148 and shall have one of the following values defined by manifest constants in the header
- 3149 <**trace.h**>:
 - 3150 `POSIX_TRACE_FLUSHING`
 - 3151 The trace stream is currently being flushed to the trace log.
 - 3152 `POSIX_TRACE_NOT_FLUSHING`
 - 3153 No flush operation is in progress.
- 3154 The *posix_stream_flush_status* member shall be set to `POSIX_TRACE_FLUSHING` if a flush
- 3155 operation is in progress either due to a call to the *posix_trace_flush()* function (explicit or caused
- 3156 by a trace stream shutdown operation) or because the trace stream has become full with the
- 3157 *stream-full-policy* attribute set to `POSIX_TRACE_FLUSH`. The *posix_stream_flush_status* member
- 3158 shall be set to `POSIX_TRACE_NOT_FLUSHING` if no flush operation is in progress.
- 3159 The *posix_stream_flush_error* member shall be set to zero if no error occurred during flushing. If
- 3160 an error occurred during a previous flushing operation, the *posix_stream_flush_error* member
- 3161 shall be set to the value of the first error that occurred. If more than one error occurs while
- 3162 flushing, error values after the first shall be discarded. The *posix_stream_flush_error* member is
- 3163 reset to zero after its value is read.
- 3164 The *posix_log_overrun_status* member indicates whether trace events were lost in the trace log,
- 3165 and shall have one of the following values defined by manifest constants in the <**trace.h**>
- 3166 header:
 - 3167 `POSIX_TRACE_OVERRUN`
 - 3168 At least one trace event was lost.
 - 3169 `POSIX_TRACE_NO_OVERRUN`
 - 3170 No trace events were lost.
- 3171 When the corresponding trace stream is created, the *posix_log_overrun_status* member shall be set
- 3172 to `POSIX_TRACE_NO_OVERRUN`. Whenever an overrun occurs, this status shall be set to
- 3173 `POSIX_TRACE_OVERRUN`. The *posix_log_overrun_status* member is reset to zero after its value
- 3174 is read.
- 3175 The *posix_log_full_status* member indicates the full status of the trace log, and it shall have one of
- 3176 the following values defined by manifest constants in the <**trace.h**> header:
 - 3177 `POSIX_TRACE_FULL`
 - 3178 The space in the trace log is exhausted.
 - 3179 `POSIX_TRACE_NOT_FULL`
 - 3180 There is still space available in the trace log.

3181 The *posix_log_full_status* member is only meaningful if the *log-full-policy* attribute is either
3182 POSIX_TRACE_UNTIL_FULL or POSIX_TRACE_LOOP.

3183 For an active trace stream without log, that is created by the *posix_trace_create()* function, the
3184 *posix_log_overrun_status* member shall be set to POSIX_TRACE_NO_OVERRUN and the
3185 *posix_log_full_status* member shall be set to POSIX_TRACE_NOT_FULL.

3186 **posix_trace_event_info Structure**

3187 The trace event structure **posix_trace_event_info** contains the information for one recorded
3188 trace event. This structure is returned by the set of functions *posix_trace_getnext_event()*,
3189 *posix_trace_timedgetnext_event()*, and *posix_trace_trygetnext_event()*.

3190 The **posix_trace_event_info** structure defined in <trace.h> shall contain at least the following
3191 members:

3192

3193

Member Type	Member Name	Description
trace_event_id_t	<i>posix_event_id</i>	Trace event type identification.
pid_t	<i>posix_pid</i>	Process ID of the process that generated the trace event.
void*	<i>posix_prog_address</i>	Address at which the trace point was invoked.
int	<i>posix_truncation_status</i>	Status about the truncation of the data associated with this trace event.
struct timespec	<i>posix_timestamp</i>	Time at which the trace event was generated.

3200

3201 In addition, if the Trace option and the Threads option are both supported, the
3202 **posix_trace_event_info** structure defined in <trace.h> shall contain the following additional
3203 member:

3204

3205

Member Type	Member Name	Description
pthread_t	<i>posix_thread_id</i>	Thread ID of the thread that generated the trace event.

3206

3207

3208

3209 The *posix_event_id* member represents the identification of the trace event type and its value is
3210 not directly defined by the user. This identification is returned by a call to one of the following
3211 functions: *posix_trace_trid_eventid_open()*, *posix_trace_eventtypelist_getnext_id()*, or
3212 *posix_trace_eventid_open()*. The name of the trace event type can be obtained by calling
3213 *posix_trace_eventid_get_name()*.

3214 The *posix_pid* is the process identifier of the traced process which generated the trace event. If
3215 the *posix_event_id* member is one of the implementation-defined system trace events and that
3216 trace event is not associated with any process, the *posix_pid* member shall be set to zero.

3217 For a user trace event, the *posix_prog_address* member is the process mapped address of the point
3218 at which the associated call to the *posix_trace_event()* function was made. For a system trace
3219 event, if the trace event is caused by a system service explicitly called by the application, the
3220 *posix_prog_address* member shall be the address of the process at the point where the call to that
3221 system service was made.

3222 The *posix_truncation_status* member defines whether the data associated with a trace event has
3223 been truncated at the time the trace event was generated, or at the time the trace event was read
3224 from the trace stream, or (if the Trace Log option is supported) from the trace log (see the *event*
3225 argument from the *posix_trace_getnext_event()* function). The *posix_truncation_status* member

3226 shall have one of the following values defined by manifest constants in the `<trace.h>` header:

3227 `POSIX_TRACE_NOT_TRUNCATED`

3228 All the traced data is available.

3229 `POSIX_TRACE_TRUNCATED_RECORD`

3230 Data was truncated at the time the trace event was generated.

3231 `POSIX_TRACE_TRUNCATED_READ`

3232 Data was truncated at the time the trace event was read from a trace stream or a trace log
3233 because the reader's buffer was too small. This truncation status overrides the
3234 `POSIX_TRACE_TRUNCATED_RECORD` status.

3235 The *posix_timestamp* member shall be the time at which the trace event was generated. The clock
3236 used is implementation-defined, but the resolution of this clock can be retrieved by a call to the
3237 *posix_trace_attr_getclockres()* function.

3238 If the Threads option is supported in addition to the Trace option:

- 3239 • The *posix_thread_id* member is the identifier of the thread that generated the trace event. If
- 3240 the *posix_event_id* member is one of the implementation-defined system trace events and that
- 3241 trace event is not associated with any thread, the *posix_thread_id* member shall be set to zero.

3242 Otherwise, this member is undefined.

3243 2.11.1.2 Trace Stream Attributes

3244 Trace streams have attributes that compose the `posix_trace_attr_t` trace stream attributes object.
3245 This object shall contain at least the following attributes:

- 3246 • The *generation-version* attribute identifies the origin and version of the trace system.
- 3247 • The *trace-name* attribute is a character string defined by the trace controller, and that
- 3248 identifies the trace stream.
- 3249 • The *creation-time* attribute represents the time of the creation of the trace stream.
- 3250 • The *clock-resolution* attribute defines the clock resolution of the clock used to generate
- 3251 timestamps.
- 3252 • The *stream-min-size* attribute defines the minimum size in bytes of the trace stream strictly
- 3253 reserved for the trace events.
- 3254 • The *stream-full-policy* attribute defines the policy followed when the trace stream is full; its
- 3255 value is `POSIX_TRACE_LOOP`, `POSIX_TRACE_UNTIL_FULL`, or `POSIX_TRACE_FLUSH`.
- 3256 • The *max-data-size* attribute defines the maximum record size in bytes of a trace event.

3257 In addition, if the Trace option and the Trace Inherit option are both supported, the
3258 `posix_trace_attr_t` trace stream creation attributes object shall contain at least the following
3259 attributes:

- 3260 • The *inheritance* attribute specifies whether a newly created trace stream will inherit tracing in
- 3261 its parent's process trace stream. It is either `POSIX_TRACE_INHERITED` or
- 3262 `POSIX_TRACE_CLOSE_FOR_CHILD`.

3263 In addition, if the Trace option and the Trace Log option are both supported, the
3264 `posix_trace_attr_t` trace stream creation attributes object shall contain at least the following
3265 attribute:

- 3266 • If the file type corresponding to the trace log supports the `POSIX_TRACE_LOOP` or the
- 3267 `POSIX_TRACE_UNTIL_FULL` policies, the *log-max-size* attribute defines the maximum size

- 3268 in bytes of the trace log associated with an active trace stream. Other stream data—for
3269 example, trace attribute values—shall not be included in this size.
- 3270 • The *log-full-policy* attribute defines the policy of a trace log associated with an active trace
3271 stream to be POSIX_TRACE_LOOP, POSIX_TRACE_UNTIL_FULL, or
3272 POSIX_TRACE_APPEND.

3273 2.11.2 Trace Event Type Definitions

3274 2.11.2.1 System Trace Event Type Definitions

3275 The following system trace event types, defined in the `<trace.h>` header, track the invocation of
3276 the trace operations:

- 3277 • POSIX_TRACE_START shall be associated with a trace start operation.
- 3278 • POSIX_TRACE_STOP shall be associated with a trace stop operation.
- 3279 • if the Trace Event Filter option is supported, POSIX_TRACE_FILTER shall be associated with
3280 a trace event type filter change operation.

3281 The following system trace event types, defined in the `<trace.h>` header, report operational trace
3282 events:

- 3283 • POSIX_TRACE_OVERFLOW shall mark the beginning of a trace overflow condition.
- 3284 • POSIX_TRACE_RESUME shall mark the end of a trace overflow condition.
- 3285 • If the Trace Log option is supported, POSIX_TRACE_FLUSH_START shall mark the
3286 beginning of a flush operation.
- 3287 • If the Trace Log option is supported, POSIX_TRACE_FLUSH_STOP shall mark the end of a
3288 flush operation.
- 3289 • If an implementation-defined trace error condition is reported, it shall be marked
3290 POSIX_TRACE_ERROR.

3291 The interpretation of a trace stream or a trace log by a trace analyzer process relies on the
3292 information recorded for each trace event, and also on system trace events that indicate the
3293 invocation of trace control operations and trace system operational trace events.

3294 The POSIX_TRACE_START and POSIX_TRACE_STOP trace events specify the time windows
3295 during which the trace stream is running.

3296 The POSIX_TRACE_STOP trace event with an associated data that is equal to zero indicates
3297 a call of the function `posix_trace_stop()`.

3298 The POSIX_TRACE_STOP trace event with an associated data that is different from zero
3299 indicates an automatic stop of the trace stream (see `posix_trace_attr_getstreamfullpolicy()`
3300 defined in the System Interfaces volume of IEEE Std. 1003.1-200x).

3301 The POSIX_TRACE_FILTER trace event indicates that a trace event type filter value changed
3302 while the trace stream was running.

3303 The POSIX_TRACE_ERROR serves to inform the analyzer process that an implementation-
3304 defined internal error of the trace system occurred.

3305 The POSIX_TRACE_OVERFLOW trace event shall be reported with a timestamp equal to the
3306 timestamp of the first trace event overwritten. This is an indication that some generated trace
3307 events have been lost.

3308 The POSIX_TRACE_RESUME trace event shall be reported with a timestamp equal to the
 3309 timestamp of the first valid trace event reported after the overflow condition ends and shall be
 3310 reported before this first valid trace event. This is an indication that the trace system is reliably
 3311 recording trace events after an overflow condition.

3312 Each of these trace event types is defined by a constant trace event name and a **trace_event_id_t**
 3313 constant; trace event data is associated with some of these trace events.

3314 If the Trace option is supported and the Trace Event Filter option and the Trace Log option are
 3315 not supported, the following predefined system trace events in Table 2-6 shall be defined:

3316 **Table 2-6** Trace Option: System Trace Events

Event Name	Constant	Associated Data
		Data Type
"posix_trace_error"	POSIX_TRACE_ERROR	error int
"posix_trace_start"	POSIX_TRACE_START	None.
"posix_trace_stop"	POSIX_TRACE_STOP	auto int
"posix_trace_overflow"	POSIX_TRACE_OVERFLOW	None.
"posix_trace_resume"	POSIX_TRACE_RESUME	None.

3326 If the Trace option and the Trace Event Filter option are both supported, and if the Trace Log
 3327 option is not supported, the following predefined system trace events in Table 2-7 shall be
 3328 defined:

3329 **Table 2-7** Trace and Trace Event Filter Options: System Trace Events

Event Name	Constant	Associated Data
		Data Type
"posix_trace_error"	POSIX_TRACE_ERROR	error int
"posix_trace_start"	POSIX_TRACE_START	event_filter trace_event_set_t
"posix_trace_stop"	POSIX_TRACE_STOP	auto int
"posix_trace_filter"	POSIX_TRACE_FILTER	old_event_filter new_event_filter trace_event_set_t
"posix_trace_overflow"	POSIX_TRACE_OVERFLOW	None.
"posix_trace_resume"	POSIX_TRACE_RESUME	None.

3343 If the Trace option and the Trace Log option are both supported, and if the Trace Event Filter
 3344 option is not supported, the following predefined system trace events in Table 2-8 (on page 581)
 3345 shall be defined:

3346

Table 2-8 Trace and Trace Log Options: System Trace Events

3347

3348

3349

3350

3351

3352

3353

3354

3355

3356

3357

3358

Event Name	Constant	Associated Data
		Data Type
"posix_trace_error"	POSIX_TRACE_ERROR	error - int
"posix_trace_start"	POSIX_TRACE_START	None.
"posix_trace_stop"	POSIX_TRACE_STOP	auto
		int
"posix_trace_overflow"	POSIX_TRACE_OVERFLOW	None.
"posix_trace_resume"	POSIX_TRACE_RESUME	None.
"posix_trace_flush_start"	POSIX_TRACE_FLUSH_START	None.
"posix_trace_flush_stop"	POSIX_TRACE_FLUSH_STOP	None.

3359

3360

If the Trace option, the Trace Event Filter option, and the Trace Log option are all supported, the following predefined system trace events in Table 2-9 shall be defined:

3361

Table 2-9 Trace, Trace Log, and Trace Event Filter Options: System Trace Events

3362

3363

3364

3365

3366

3367

3368

3369

3370

3371

3372

3373

3374

3375

3376

Event Name	Constant	Associated Data
		Data Type
"posix_trace_error"	POSIX_TRACE_ERROR	error
		int
"posix_trace_start"	POSIX_TRACE_START	event_filter
		trace_event_set_t
"posix_trace_stop"	POSIX_TRACE_STOP	auto
		int
"posix_trace_filter"	POSIX_TRACE_FILTER	old_event_filter
		new_event_filter
		trace_event_set_t
"posix_trace_overflow"	POSIX_TRACE_OVERFLOW	None.
"posix_trace_resume"	POSIX_TRACE_RESUME	None.
"posix_trace_flush_start"	POSIX_TRACE_FLUSH_START	None.
"posix_trace_flush_stop"	POSIX_TRACE_FLUSH_STOP	None.

3377 2.11.2.2 User Trace Event Type Definitions

3378

3379

3380

3381

3382

The user trace event POSIX_TRACE_UNNAMED_USEREVENT shall be defined in the `<trace.h>` header. If the limit of per-process user trace event names represented by {TRACE_USER_EVENT_MAX} has already been reached, this predefined user event shall be returned when the application tries to register more events than allowed. The data associated with this trace event is application-defined.

3383

The following predefined user trace event in Table 2-10 (on page 582) shall be defined:

3384

Table 2-10 Trace Option: User Trace Event

3385

3386

Event Name	Constant
"posix_trace_unnamed_userevent"	POSIX_TRACE_UNNAMED_USEREVENT

3387 **2.11.3 Trace Functions**

3388

3389

3390

3391

3392

3393

The trace interface is built and structured to improve portability through use of trace data of opaque type. The object-oriented approach for the manipulation of trace attributes and trace event type identifiers requires definition of many constructor and selector functions which operate on these opaque types. Also, the trace interface must support several different tracing roles. To facilitate reading the trace interface, the trace functions are grouped into small functional sets supporting the three different roles:

3394

3395

- A trace controller process requires functions to set up and customize all the resources needed to run a trace stream, including:

3396

- Attribute initialization and destruction (*posix_trace_attr_init()*)

3397

- Identification information manipulation (*posix_trace_attr_getgenversion()*)

3398

- Trace system behavior modification (*posix_trace_attr_getinherited()*)

3399

- Trace stream and trace log size set (*posix_trace_attr_getmaxusersize()*)

3400

- Trace stream creation, flush, and shutdown (*posix_trace_create()*)

3401

- Trace stream and trace log clear (*posix_trace_clear()*)

3402

- Trace event type identifier manipulation (*posix_trace_trid_eventid_open()*)

3403

- Trace event type identifier list exploration (*posix_trace_eventtypelist_getnext_id()*)

3404

- Trace event type set manipulation (*posix_trace_eventset_empty()*)

3405

- Trace event type filter set (*posix_trace_set_filter()*)

3406

- Trace stream start and stop (*posix_trace_start()*)

3407

- Trace stream information and status read (*posix_trace_get_attr()*)

3408

- A traced process requires functions to instrument trace points:

3409

- Trace event type identifiers definition and trace points insertion (*posix_trace_event()*)

3410

3411

- A trace analyzer process requires functions to retrieve information from a trace stream and trace log:

3412

- Identification information read (*posix_trace_attr_getgenversion()*)

3413

- Trace system behavior information read (*posix_trace_attr_getinherited()*)

3414

- Trace stream and trace log size get (*posix_trace_attr_getmaxusersize()*)

3415

- Trace event type identifier manipulation (*posix_trace_trid_eventid_open()*)

3416

- Trace event type identifier list exploration (*posix_trace_eventtypelist_getnext_id()*)

3417

- Trace log open, rewind, and close (*posix_trace_open()*)

3418

- Trace stream information and status read (*posix_trace_get_attr()*)

3419

- trace event read (*posix_trace_getnext_event()*)

3420 **2.12 Data Types**

3421 All of the data types used by various functions are defined by the implementation. The
 3422 following table describes some of these types. Other types referenced in the description of a
 3423 function, not mentioned here, can be found in the appropriate header for that function.

3424

3425

Defined Type	Description
3426 cc_t	Type used for terminal special characters.
3427 clock_t	Arithmetic type used for processor times, as defined in the ISO C standard.
3428	
3429 clockid_t	Used for clock ID type in some timer functions.
3430 dev_t	Arithmetic type used for device numbers.
3431 DIR	Type representing a directory stream.
3432 div_t	Structure type returned by the <i>div()</i> function.
3433 FILE	Structure containing information about a file.
3434 glob_t	Structure type used in path name pattern matching.
3435 fpos_t	Type containing all information needed to specify uniquely every position within a file.
3436	
3437 gid_t	Arithmetic type used for group IDs.
3438 iconv_t	Type used for conversion descriptors.
3439 id_t	Arithmetic type used as a general identifier; can be used to contain at least the largest of a pid_t , uid_t , or gid_t .
3440	
3441 ino_t	Arithmetic type used for file serial numbers.
3442 key_t	Arithmetic type used for XSI interprocess communication.
3443 ldiv_t	Structure type returned by the <i>ldiv()</i> function.
3444 mode_t	Arithmetic type used for file attributes.
3445 mqd_t	Used for message queue descriptors.
3446 nfds_t	Integer type used for the number of file descriptors.
3447 nlink_t	Arithmetic type used for link counts.
3448 off_t	Signed arithmetic type used for file sizes.
3449 pid_t	Signed arithmetic type used for process and process group IDs.
3450 pthread_attr_t	Used to identify a thread attribute object.
3451 pthread_cond_t	Used for condition variables.
3452 pthread_condattr_t	Used to identify a condition attribute object.
3453 pthread_key_t	Used for thread-specific data keys.
3454 pthread_mutex_t	Used for mutexes.
3455 pthread_mutexattr_t	Used to identify a mutex attribute object.
3456 pthread_once_t	Used for dynamic package initialization.
3457 pthread_rwlock_t	Used for read-write locks.
3458 pthread_rwlockattr_t	Used for read-write lock attributes.
3459 pthread_t	Used to identify a thread.
3460 ptrdiff_t	Signed integer type of the result of subtracting two pointers.
3461 regex_t	Structure type used in regular expression matching.
3462 regmatch_t	Structure type used in regular expression matching.
3463 rlim_t	Unsigned arithmetic type used for limit values, to which objects of type int and off_t can be cast without loss of value.
3464	
3465 sem_t	Type used in performing semaphore operations.
3466 sig_atomic_t	Integer type of an object that can be accessed as an atomic

3467
 3468
 3469
 3470
 3471
 3472
 3473
 3474
 3475
 3476
 3477
 3478
 3479
 3480
 3481
 3482
 3483
 3484
 3485
 3486
 3487
 3488
 3489
 3490

Defined Type	Description
sigset_t	entity, even in the presence of asynchronous interrupts. Integer or structure type of an object used to represent sets of signals.
size_t	Unsigned integer type used for size of objects.
speed_t	Type used for terminal baud rates.
ssize_t	Arithmetic type used for a count of bytes or an error indication.
suseconds_t	Signed arithmetic type used for time in microseconds.
tcflag_t	Type used for terminal modes.
time_t	Arithmetic type used for time in seconds, as defined in the ISO C standard.
timer_t	Used for timer ID returned by the <i>timer_create()</i> function.
uid_t	Arithmetic type used for user IDs.
useconds_t	Integer type used for time in microseconds.
va_list	Type used for traversing variable argument lists.
wchar_t	Integer type whose range of values can represent distinct codes for all members of the largest extended character set specified by the supported locales.
wctype_t	Scalar type which represents a character class descriptor.
wint_t	Integer type capable of storing any valid value of wchar_t or WEOF.
wordexp_t	Structure type used in word expansion.

System Interfaces

3491

3492 This chapter describes the functions, macros, and external variables to support applications
3493 portability at the C-language source level.

3494 **NAME**

3495 FD_CLR — macros for synchronous I/O multiplexing

3496 **SYNOPSIS**

3497 #include <sys/time.h>

3498 FD_CLR(int *fd*, fd_set **fdset*);3499 FD_ISSET(int *fd*, fd_set **fdset*);3500 FD_SET(int *fd*, fd_set **fdset*);3501 FD_ZERO(fd_set **fdset*);3502 **DESCRIPTION**3503 Refer to *select()*.

3504 **NAME**
3505 `_Exit, _exit` — terminate a process

3506 **SYNOPSIS**
3507 `#include <unistd.h>`

3508 `void _Exit(int status);`
3509 `void _exit(int status);`

3510 **DESCRIPTION**
3511 Refer to *exit()*.

3512 **NAME**

3513 _longjmp, _setjmp — non-local goto

3514 **SYNOPSIS**

3515 xSI #include <setjmp.h>

3516 void _longjmp(jmp_buf env, int val);

3517 int _setjmp(jmp_buf env);

3518

3519 **DESCRIPTION**

3520 The *_longjmp()* and *_setjmp()* functions are identical to *longjmp()* and *setjmp()*, respectively, with
3521 the additional restriction that *_longjmp()* and *_setjmp()* do not manipulate the signal mask.

3522 If *_longjmp()* is called even though *env* was never initialized by a call to *_setjmp()*, or when the
3523 last such call was in a function that has since returned, the results are undefined.

3524 **RETURN VALUE**

3525 Refer to *longjmp()* and *setjmp()*.

3526 **ERRORS**

3527 No errors are defined.

3528 **EXAMPLES**

3529 None.

3530 **APPLICATION USAGE**

3531 If *_longjmp()* is executed and the environment in which *_setjmp()* was executed no longer exists,
3532 errors can occur. The conditions under which the environment of the *_setjmp()* no longer exists
3533 include exiting the function that contains the *_setjmp()* call, and exiting an inner block with
3534 temporary storage. This condition might not be detectable, in which case the *_longjmp()* occurs
3535 and, if the environment no longer exists, the contents of the temporary storage of an inner block
3536 are unpredictable. This condition might also cause unexpected process termination. If the
3537 function has returned, the results are undefined.

3538 Passing *longjmp()* a pointer to a buffer not created by *setjmp()*, passing *_longjmp()* a pointer to a
3539 buffer not created by *_setjmp()*, passing *siglongjmp()* a pointer to a buffer not created by
3540 *sigsetjmp()*, or passing any of these three functions a buffer that has been modified by the user
3541 can cause all the problems listed above, and more.

3542 The *_longjmp()* and *_setjmp()* functions are included to support programs written to historical
3543 system interfaces. New applications should use *siglongjmp()* and *sigsetjmp()* respectively.

3544 **RATIONALE**

3545 None.

3546 **FUTURE DIRECTIONS**

3547 The *_longjmp()* and *_setjmp()* functions may be marked LEGACY in a future version.

3548 **SEE ALSO**

3549 *longjmp()*, *setjmp()*, *siglongjmp()*, *sigsetjmp()*, the Base Definitions volume of
3550 IEEE Std. 1003.1-200x, <setjmp.h>

3551 **CHANGE HISTORY**

3552 First released in Issue 4, Version 2.

3553 **Issue 5**

3554 Moved from X/OPEN UNIX extension to BASE.

3555 **NAME**

3556 _setjmp — set jump point for a non-local goto

3557 **SYNOPSIS**

3558 xSI #include <setjmp.h>

3559 int _setjmp(jmp_buf env);

3560

3561 **DESCRIPTION**

3562 Refer to *_longjmp()*.

3563 **NAME**

3564 _toupper — transliterate uppercase characters to lowercase

3565 **SYNOPSIS**

3566 XSI #include <ctype.h>

3567 int _tolower(int c);

3568

3569 **DESCRIPTION**3570 The *_tolower()* macro shall be equivalent to *tolower(c)* except that the application shall ensure
3571 that the argument *c* is an uppercase letter.3572 **RETURN VALUE**3573 Upon successful completion, *_tolower()* shall return the lowercase letter corresponding to the
3574 argument passed.3575 **ERRORS**

3576 No errors are defined.

3577 **EXAMPLES**

3578 None.

3579 **APPLICATION USAGE**

3580 None.

3581 **RATIONALE**

3582 None.

3583 **FUTURE DIRECTIONS**

3584 None.

3585 **SEE ALSO**3586 *tolower()*, *isupper()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base
3587 Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale3588 **CHANGE HISTORY**

3589 First released in Issue 1. Derived from Issue 1 of the SVID.

3590 **Issue 4**

3591 The RETURN VALUE section is expanded.

3592 **Issue 6**

3593 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

3594 **NAME**

3595 _toupper — transliterate lowercase characters to uppercase

3596 **SYNOPSIS**

3597 xSI #include <ctype.h>

3598 int _toupper(int c);

3599

3600 **DESCRIPTION**

3601 The *_toupper()* macro shall be equivalent to *toupper()* except that the application shall ensure
3602 that the argument *c* is a lowercase letter.

3603 **RETURN VALUE**

3604 Upon successful completion, *_toupper()* shall return the uppercase letter corresponding to the
3605 argument passed.

3606 **ERRORS**

3607 No errors are defined.

3608 **EXAMPLES**

3609 None.

3610 **APPLICATION USAGE**

3611 None.

3612 **RATIONALE**

3613 None.

3614 **FUTURE DIRECTIONS**

3615 None.

3616 **SEE ALSO**

3617 *islower()*, *toupper()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base
3618 Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

3619 **CHANGE HISTORY**

3620 First released in Issue 1. Derived from Issue 1 of the SVID.

3621 **Issue 4**

3622 The RETURN VALUE section is expanded.

3623 **Issue 6**

3624 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

3625 **NAME**

3626 a64l, l64a — convert between a 32-bit integer and a radix-64 ASCII string

3627 **SYNOPSIS**

```
3628 xSI      #include <stdlib.h>
3629
3629         long a64l(const char *s);
3630         char *l64a(long value);
3631
```

3632 **DESCRIPTION**

3633 These functions are used to maintain numbers stored in radix-64 ASCII characters. This is a
 3634 notation by which 32-bit integers can be represented by up to six characters; each character
 3635 represents a digit in radix-64 notation. If the type **long** contains more than 32 bits, only the low-
 3636 order 32 bits shall be used for these operations.

3637 The characters used to represent digits are '.' (dot) for 0, '/' for 1, '0' through '9' for 2-11,
 3638 'A' through 'Z' for 12-37, and 'a' through 'z' for 38-63.

3639 The *a64l()* function shall take a pointer to a radix-64 representation, in which the first digit is the
 3640 least significant, and return a corresponding **long** value. If the string pointed to by *s* contains
 3641 more than six characters, *a64l()* shall use the first six. If the first six characters of the string
 3642 contain a null terminator, *a64l()* shall use only characters preceding the null terminator. The
 3643 *a64l()* function scans the character string from left to right with the least significant digit on the
 3644 left, decoding each character as a 6-bit radix-64 number. If the type **long** contains more than 32
 3645 bits, the resulting value is sign-extended. The behavior of *a64l()* is unspecified if *s* is a null
 3646 pointer or the string pointed to by *s* was not generated by a previous call to *l64a()*.

3647 The *l64a()* function shall take a **long** argument and return a pointer to the corresponding radix-
 3648 64 representation. The behavior of *l64a()* is unspecified if *value* is negative.

3649 The value returned by *l64a()* may be a pointer into a static buffer. Subsequent calls to *l64a()* may
 3650 overwrite the buffer.

3651 The *l64a()* function need not be reentrant. A function that is not required to be reentrant is not
 3652 required to be thread-safe.

3653 **RETURN VALUE**

3654 Upon successful completion, *a64l()* shall return the **long** value resulting from conversion of the
 3655 input string. If a string pointed to by *s* is an empty string, *a64l()* shall return 0L.

3656 The *l64a()* function shall return a pointer to the radix-64 representation. If *value* is 0L, *l64a()* shall
 3657 return a pointer to an empty string.

3658 **ERRORS**

3659 No errors are defined.

3660 **EXAMPLES**

3661 None.

3662 **APPLICATION USAGE**

3663 If the type **long** contains more than 32 bits, the result of *a64l(l64a(x))* is *x* in the low-order 32 bits.

3664 **RATIONALE**

3665 This is not the same encoding as used by either encoding variant of the *uuencode* utility.

3666 **FUTURE DIRECTIONS**

3667 None.

3668 **SEE ALSO**3669 *strtoul()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>`, the Shell and Utilities
3670 volume of IEEE Std. 1003.1-200x, *uuencode*3671 **CHANGE HISTORY**

3672 First released in Issue 4, Version 2.

3673 **Issue 5**

3674 Moved from X/OPEN UNIX extension to BASE.

3675 Normative text previously in the APPLICATION USAGE section moved to the DESCRIPTION.

3676 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

3677 **NAME**

3678 abort — generate an abnormal process abort

3679 **SYNOPSIS**

3680 #include <stdlib.h>

3681 void abort(void);

3682 **DESCRIPTION**

3683 CX The functionality described on this reference page is aligned with the ISO C standard. Any
3684 conflict between the requirements described here and the ISO C standard is unintentional. This
3685 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

3686 The *abort()* function shall cause abnormal process termination to occur, unless the signal
3687 SIGABRT is being caught and the signal handler does not return.

3688 CX The abnormal termination processing shall include at least the effect of *fclose()* on all open
3689 streams and the default actions defined for SIGABRT.

3690 XSI On XSI-conformant systems, in addition the abnormal termination processing shall include the
3691 effect of *fclose()* on message catalog descriptors.

3692 The SIGABRT signal shall be sent to the calling process as if by means of *raise()* with the
3693 argument SIGABRT.

3694 CX The status made available to *wait()* or *waitpid()* by *abort()* shall be that of a process terminated
3695 by the SIGABRT signal. The *abort()* function shall override blocking or ignoring the SIGABRT
3696 signal.

3697 **RETURN VALUE**3698 The *abort()* function shall not return.3699 **ERRORS**

3700 No errors are defined.

3701 **EXAMPLES**

3702 None.

3703 **APPLICATION USAGE**

3704 Catching the signal is intended to provide the application writer with a portable means to abort
3705 processing, free from possible interference from any implementation-defined library functions. If
3706 SIGABRT is neither caught nor ignored, then the actions associated with SIGABRT defined in
3707 Section 2.4.1 (on page 528) will be taken.

3708 **RATIONALE**

3709 None.

3710 **FUTURE DIRECTIONS**

3711 None.

3712 **SEE ALSO**

3713 *exit()*, *kill()*, *raise()*, *signal()*, *wait()*, *waitpid()*, the Base Definitions volume of
3714 IEEE Std. 1003.1-200x, <stdlib.h>

3715 **CHANGE HISTORY**

3716 First released in Issue 1. Derived from Issue 1 of the SVID.

3717 Issue 4

3718 The following changes are incorporated in this issue for alignment with the ISO C standard and
3719 the ISO POSIX-1 standard:

- 3720 • The argument list is explicitly defined as **void**.
- 3721 • The DESCRIPTION is revised to identify the correct order in which operations occur. It also
3722 identifies:
 - 3723 — How the calling process is signaled
 - 3724 — How status information is made available to the host environment
 - 3725 — That *abort()* overrides blocking or ignoring of the SIGABRT signal
- 3726 • The APPLICATION USAGE section is replaced.

3727 Issue 6

3728 Extensions beyond the ISO C standard are now marked.

3729 **NAME**

3730 abs — return an integer absolute value

3731 **SYNOPSIS**

3732 #include <stdlib.h>

3733 int abs(int *i*);

3734 **DESCRIPTION**

3735 cx The functionality described on this reference page is aligned with the ISO C standard. Any
3736 conflict between the requirements described here and the ISO C standard is unintentional. This
3737 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

3738 The *abs()* function shall compute the absolute value of its integer operand, *i*. If the result cannot
3739 be represented, the behavior is undefined.

3740 **RETURN VALUE**

3741 The *abs()* function shall return the absolute value of its integer operand.

3742 **ERRORS**

3743 No errors are defined.

3744 **EXAMPLES**

3745 None.

3746 **APPLICATION USAGE**

3747 In two's-complement representation, the absolute value of the negative integer with largest
3748 magnitude {INT_MIN} might not be representable.

3749 **RATIONALE**

3750 None.

3751 **FUTURE DIRECTIONS**

3752 None.

3753 **SEE ALSO**

3754 *fabs()*, *labs()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

3755 **CHANGE HISTORY**

3756 First released in Issue 1. Derived from Issue 1 of the SVID.

3757 **Issue 4**

3758 In the APPLICATION USAGE section, the phrase “{INT_MIN} is undefined” is replaced with
3759 “{INT_MIN} might not be representable”.

3760 **Issue 6**

3761 Extensions beyond the ISO C standard are now marked.

3762 NAME

3763 accept — accept a new connection on a socket

3764 SYNOPSIS

3765 #include <sys/socket.h>

3766 int accept(int *socket*, struct sockaddr *restrict *address*,
3767 socklen_t *restrict *address_len*);

3768 DESCRIPTION

3769 The *accept()* function shall extract the first connection on the queue of pending connections,
3770 create a new socket with the same socket type protocol and address family as the specified
3771 socket, and allocate a new file descriptor for that socket.3772 The *accept()* function takes the following arguments:3773 *socket* Specifies a socket that was created with *socket()*, has been bound to an address
3774 with *bind()*, and has issued a successful call to *listen()*.3775 *address* Either a null pointer, or a pointer to a **sockaddr** structure where the address of
3776 the connecting socket shall be returned.3777 *address_len* Points to a **socklen_t** structure which on input specifies the length of the
3778 supplied **sockaddr** structure, and on output specifies the length of the stored
3779 address.3780 If *address* is not a null pointer, the address of the peer for the accepted connection shall be stored
3781 in the **sockaddr** structure pointed to by *address*, and the length of this address shall be stored in
3782 the object pointed to by *address_len*.3783 If the actual length of the address is greater than the length of the supplied **sockaddr** structure,
3784 the stored address shall be truncated.3785 If the protocol permits connections by unbound clients, and the peer is not bound, then the value
3786 stored in the object pointed to by *address* is unspecified.3787 If the listen queue is empty of connection requests and O_NONBLOCK is not set on the file
3788 descriptor for the socket, *accept()* shall block until a connection is present. If the *listen()* queue is
3789 empty of connection requests and O_NONBLOCK is set on the file descriptor for the socket,
3790 *accept()* shall fail and set *errno* to [EAGAIN] or [EWOULDBLOCK].3791 The accepted socket cannot itself accept more connections. The original socket remains open and
3792 can accept more connections.

3793 RETURN VALUE

3794 Upon successful completion, *accept()* shall return the non-negative file descriptor of the accepted
3795 socket. Otherwise, -1 shall be returned and *errno* set to indicate the error.

3796 ERRORS

3797 The *accept()* function shall fail if:

3798 [EAGAIN] or [EWOULDBLOCK]

3799 O_NONBLOCK is set for the socket file descriptor and no connections are
3800 present to be accepted.3801 [EBADF] The *socket* argument is not a valid file descriptor.

3802 [ECONNABORTED]

3803 A connection has been aborted.

3804	[EINTR]	The <i>accept()</i> function was interrupted by a signal that was caught before a valid connection arrived.
3805		
3806	[EINVAL]	The <i>socket</i> is not accepting connections.
3807	[EMFILE]	{OPEN_MAX} file descriptors are currently open in the calling process.
3808	[ENFILE]	The maximum number of file descriptors in the system are already open.
3809	[ENOTSOCK]	The <i>socket</i> argument does not refer to a socket.
3810	[EOPNOTSUPP]	The <i>socket</i> type of the specified socket does not support accepting connections.
3811		
3812		The <i>accept()</i> function may fail if:
3813	[ENOBUFS]	No buffer space is available.
3814	[ENOMEM]	There was insufficient memory available to complete the operation.
3815	XSR [EPROTO]	A protocol error has occurred; for example, the STREAMS protocol stack has not been initialized.
3816		

3817 EXAMPLES

3818 None.

3819 APPLICATION USAGE

3820 When a connection is available, *select()* indicates that the file descriptor for the socket is ready
3821 for reading.

3822 RATIONALE

3823 None.

3824 FUTURE DIRECTIONS

3825 None.

3826 SEE ALSO

3827 *bind()*, *connect()*, *listen()*, *socket()*, the Base Definitions volume of IEEE Std. 1003.1-200x,
3828 <sys/socket.h>

3829 CHANGE HISTORY

3830 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

3831 The **restrict** keyword is added to the *accept()* prototype for alignment with the
3832 ISO/IEC 9899:1999 standard.

3833 **NAME**

3834 access — determine accessibility of a file

3835 **SYNOPSIS**

3836 #include <unistd.h>

3837 int access(const char *path, int amode);

3838 **DESCRIPTION**

3839 The `access()` function shall check the file named by the `path` name pointed to by the `path`
3840 argument for accessibility according to the bit pattern contained in `amode`, using the real user ID
3841 in place of the effective user ID and the real group ID in place of the effective group ID.

3842 The value of `amode` is either the bitwise-inclusive OR of the access permissions to be checked
3843 (R_OK, W_OK, X_OK) or the existence test (F_OK).

3844 If any access permissions are checked, each shall be checked individually, as described in the
3845 Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 3, Definitions. If the process has
3846 appropriate privileges, an implementation may indicate success for X_OK even if none of the
3847 execute file permission bits are set.

3848 **RETURN VALUE**

3849 If the requested access is permitted, `access()` succeeds and shall return 0; otherwise, -1 shall be
3850 returned and `errno` shall be set to indicate the error.

3851 **ERRORS**3852 The `access()` function shall fail if:

3853 [EACCES] Permission bits of the file mode do not permit the requested access, or search
3854 permission is denied on a component of the path prefix.

3855 [ELOOP] A loop exists in symbolic links encountered during resolution of the `path`
3856 argument.

3857 [ENAMETOOLONG]
3858 The length of the `path` argument exceeds {PATH_MAX} or a path name
3859 component is longer than {NAME_MAX}.

3860 [ENOENT] A component of `path` does not name an existing file or `path` is an empty string.

3861 [ENOTDIR] A component of the path prefix is not a directory.

3862 [EROFS] Write access is requested for a file on a read-only file system.

3863 The `access()` function may fail if:

3864 [EINVAL] The value of the `amode` argument is invalid.

3865 [ELOOP] More than {SYMLOOP_MAX} symbolic links were encountered during
3866 resolution of the `path` argument.

3867 [ENAMETOOLONG]
3868 As a result of encountering a symbolic link in resolution of the `path` argument,
3869 the length of the substituted path name string exceeded {PATH_MAX}.

3870 [ETXTBSY] Write access is requested for a pure procedure (shared text) file that is being
3871 executed.

3872 **EXAMPLES**3873 **Testing for the Existence of a File**

3874 The following example tests whether a file named **myfile** exists in the **/tmp** directory.

```
3875 #include <unistd.h>
3876 ...
3877 int result;
3878 const char *filename = "/tmp/myfile";
3879 result = access (filename, F_OK);
```

3880 **APPLICATION USAGE**

3881 Additional values of *amode* other than the set defined in the description may be valid; for
3882 example, if a system has extended access controls.

3883 **RATIONALE**

3884 In early proposals, some inadequacies in the *access()* function led to the creation of an *eaccess()*
3885 function because:

- 3886 1. Historical implementations of *access()* do not test file access correctly when the process'
3887 real user ID is superuser. In particular, they always return zero when testing execute
3888 permissions without regard to whether the file is executable.
- 3889 2. The superuser has complete access to all files on a system. As a consequence, programs
3890 started by the superuser and switched to the effective user ID with lesser privileges cannot
3891 use *access()* to test their file access permissions.

3892 However, the historical model of *eaccess()* does not resolve problem (1), so this volume of
3893 IEEE Std. 1003.1-200x now allows *access()* to behave in the desired way because several
3894 implementations have corrected the problem. It was also argued that problem (2) is more easily
3895 solved by using *open()*, *chdir()*, or one of the *exec* functions as appropriate and responding to the
3896 error, rather than creating a new function that would not be as reliable. Therefore, *eaccess()* is not
3897 included in this volume of IEEE Std. 1003.1-200x.

3898 The sentence concerning appropriate privileges and execute permission bits reflects the two
3899 possibilities implemented by historical implementations when checking superuser access for
3900 *X_OK*.

3901 **FUTURE DIRECTIONS**

3902 None.

3903 **SEE ALSO**

3904 *chmod()*, *stat()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<unistd.h>**

3905 **CHANGE HISTORY**

3906 First released in Issue 1. Derived from Issue 1 of the SVID.

3907 **Issue 4**

3908 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 3909 • The type of argument *path* is changed from **char*** to **const char***.

3910 The following change is incorporated for alignment with the FIPS requirements:

- 3911 • In the **ERRORS** section, the condition whereby [ENAMETOOLONG] is returned if a path
3912 name component is larger than {NAME_MAX} is now defined as mandatory and marked as
3913 an extension.

3914 **Issue 4, Version 2**

3915 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 3916 • It states that [ELOOP] is returned if too many symbolic links are encountered during path
- 3917 name resolution.
- 3918 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an
- 3919 intermediate result of path name resolution of a symbolic link.

3920 **Issue 6**

3921 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 3922 • The [ENAMETOOLONG] error is restored as an error dependent on _POSIX_NO_TRUNC.
- 3923 This is since behavior may vary from one file system to another.

3924 The following new requirements on POSIX implementations derive from alignment with the
3925 Single UNIX Specification:

- 3926 • The [ELOOP] mandatory error condition is added.
- 3927 • A second [ENAMETOOLONG] is added as an optional error condition.
- 3928 • The [ETXTBSY] optional error condition is added.

3929 The following changes were made to align with the IEEE P1003.1a draft standard:

- 3930 • The [ELOOP] optional error condition is added.

3931 **NAME**

3932 acos, acosf, acosl — arc cosine function

3933 **SYNOPSIS**

3934 #include <math.h>

3935 double acos(double x);

3936 float acosf(float x);

3937 long double acosl(long double x);

3938 **DESCRIPTION**

3939 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
 3940 conflict between the requirements described here and the ISO C standard is unintentional. This
 3941 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

3942 The *acos* family of functions shall compute the principal value of the arc cosine of *x*. The value
 3943 of *x* should be in the range $[-1,1]$.

3944 An application wishing to check for error situations should set *errno* to 0 before calling *acos()*. If
 3945 *errno* is non-zero on return, or the value NaN is returned, an error has occurred.

3946 **RETURN VALUE**

3947 Upon successful completion, the *acos* family of functions shall return the arc cosine of *x*, in the
 3948 **XSI** range $[0,\pi]$ radians. If the value of *x* is not in the range $[-1,1]$, and is not $\pm\text{Inf}$ or NaN, either 0.0 or
 3949 NaN shall be returned and *errno* shall be set to [EDOM].

3950 **XSI** If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM]. If *x* is $\pm\text{Inf}$, either 0.0 shall be
 3951 returned and *errno* shall be set to [EDOM], or NaN shall be returned and *errno* may be set to
 3952 [EDOM].

3953 **ERRORS**3954 The *acos* family of functions shall fail if:

3955 **XSI** [EDOM] The value *x* is not $\pm\text{Inf}$ or NaN and is not in the range $[-1,1]$.

3956 The *acos* family of functions may fail if:

3957 **XSI** [EDOM] The value *x* is $\pm\text{Inf}$ or NaN.

3958 **XSI** No other errors shall occur.3959 **EXAMPLES**

3960 None.

3961 **APPLICATION USAGE**

3962 None.

3963 **RATIONALE**

3964 None.

3965 **FUTURE DIRECTIONS**

3966 None.

3967 **SEE ALSO**3968 *cos()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>3969 **CHANGE HISTORY**

3970 First released in Issue 1. Derived from Issue 1 of the SVID.

3971 **Issue 4**

3972 Removed references to *matherr()*.

3973 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the
3974 ISO C standard and to rationalize error handling in the mathematics functions.

3975 The return value specified for [EDOM] is marked as an extension.

3976 **Issue 5**

3977 The DESCRIPTION is updated to indicate how an application should check for an error. This
3978 text was previously published in the APPLICATION USAGE section.

3979 **Issue 6**

3980 The *acosf()* and *acosl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

3981 **NAME**3982 `acosh`, `acoshf`, `acoshl`, `asinh`, `asinhf`, `asinhf`, `atanh`, `atanhf`, `atanhl` — inverse hyperbolic functions3983 **SYNOPSIS**3984 `#include <math.h>`

```

3985     double acosh(double x);
3986     float  acoshf(float x);
3987     long double acoshl(long double x);
3988     double asinh(double x);
3989     float  asinhf(float x);
3990     long double asinhf(long double x);
3991     double atanh(double x);
3992     float  atanhf(float x);
3993     long double atanhf(long double x);

```

3994 **DESCRIPTION**

3995 The `acosh()`, `asinh()`, and `atanh()` functions shall compute the inverse hyperbolic cosine, sine, and
 3996 tangent of their argument, respectively.

3997 **RETURN VALUE**

3998 The `acosh()`, `asinh()`, and `atanh()` functions shall return the inverse hyperbolic cosine, sine, and
 3999 tangent of their argument, respectively.

4000 The `acosh()`, `acoshf()`, and `acoshl()` functions shall return an implementation-defined value (NaN
 4001 or equivalent if available) and set `errno` to [EDOM] when its argument is less than 1.0.

4002 The `atanh()`, `atanhf()`, and `atanhl()` functions shall return an implementation-defined value (NaN
 4003 or equivalent if available) and set `errno` to [EDOM] when its argument has absolute value greater
 4004 than 1.0.

4005 If `x` is NaN, the `asinh()`, `acosh()`, and `atanh()` functions shall return NaN and may set `errno` to
 4006 [EDOM].

4007 **ERRORS**

4008 The `acosh()`, `acoshf()`, and `acoshl()` functions shall fail if:

4009 [EDOM] The `x` argument is less than 1.0.

4010 The `atanh()`, `atanhf()`, and `atanhl()` functions shall fail if:

4011 [EDOM] The `x` argument has an absolute value greater than 1.0.

4012 The `atanh()`, `atanhf()`, and `atanhl()` functions shall fail if:

4013 [ERANGE] The `x` argument has an absolute value equal to 1.0

4014 The `asinh()`, `acosh()`, and `atanh()` functions may fail if:

4015 [EDOM] The value of `x` is NaN.

4016 **EXAMPLES**

4017 None.

4018 **APPLICATION USAGE**

4019 None.

4020 **RATIONALE**

4021 None.

4022 **FUTURE DIRECTIONS**

4023 None.

4024 **SEE ALSO**4025 *cosh()*, *sinh()*, *tanh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>4026 **CHANGE HISTORY**

4027 First released in Issue 4, Version 2.

4028 **Issue 5**

4029 Moved from X/OPEN UNIX extension to BASE.

4030 **Issue 6**4031 The *acoshf()*, *acoshl()*, *asinhf()*, *asinhf()*, *atanhf()*, and *atanhl()* functions are added for alignment
4032 with the ISO/IEC 9899:1999 standard.

4033 **NAME**4034 aio_cancel — cancel an asynchronous I/O request (**REALTIME**)4035 **SYNOPSIS**

4036 AIO #include <aio.h>

4037 int aio_cancel(int *fildev*, struct aiocb **aiocbp*);

4038

4039 **DESCRIPTION**

4040 The *aio_cancel()* function shall attempt to cancel one or more asynchronous I/O requests
 4041 currently outstanding against file descriptor *fildev*. The *aiocbp* argument points to the
 4042 asynchronous I/O control block for a particular request to be canceled. If *aiocbp* is NULL, then
 4043 all outstanding cancelable asynchronous I/O requests against *fildev* shall be canceled.

4044 Normal asynchronous notification shall occur for asynchronous I/O operations that are
 4045 successfully canceled. If there are requests that cannot be canceled, then the normal
 4046 asynchronous completion process shall take place for those requests when they are completed.

4047 For requested operations that are successfully canceled, the associated error status shall be set to
 4048 [ECANCELED] and the return status shall be -1. For requested operations that are not
 4049 successfully canceled, the *aiocbp* shall not be modified by *aio_cancel()*.

4050 If *aiocbp* is not NULL, then if *fildev* does not have the same value as the file descriptor with which
 4051 the asynchronous operation was initiated, unspecified results occur.

4052 Which operations are cancelable is implementation-defined.

4053 **RETURN VALUE**

4054 The *aio_cancel()* function shall return the value AIO_CANCELED to the calling process if the
 4055 requested operation(s) were canceled. The value AIO_NOTCANCELED shall be returned if at
 4056 least one of the requested operation(s) cannot be canceled because it is in progress. In this case,
 4057 the state of the other operations, if any, referenced in the call to *aio_cancel()* is not indicated by
 4058 the return value of *aio_cancel()*. The application may determine the state of affairs for these
 4059 operations by using *aio_error()*. The value AIO_ALLDONE is returned if all of the operations
 4060 have already completed. Otherwise, the function shall return -1 and set *errno* to indicate the
 4061 error.

4062 **ERRORS**

4063 The *aio_cancel()* function shall fail if:

4064 [EBADF] The *fildev* argument is not a valid file descriptor.

4065 **EXAMPLES**

4066 None.

4067 **APPLICATION USAGE**

4068 The *aio_cancel()* function is part of the Asynchronous Input and Output option and need not be
 4069 available on all implementations.

4070 **RATIONALE**

4071 None.

4072 **FUTURE DIRECTIONS**

4073 None.

4074 **SEE ALSO**4075 *aio_read()*, *aio_write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**aio.h**>4076 **CHANGE HISTORY**

4077 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4078 **Issue 6**4079 The [ENOSYS] error condition has been removed as stubs need not be provided if an
4080 implementation does not support the Asynchronous Input and Output option.

4081 The APPLICATION USAGE section is added.

4082 **NAME**4083 aio_error — retrieve errors status for an asynchronous I/O operation (**REALTIME**)4084 **SYNOPSIS**

4085 AIO #include <aio.h>

4086 int aio_error(const struct aiocb *aiocbp);

4087

4088 **DESCRIPTION**

4089 The *aio_error()* function shall return the error status associated with the **aiocb** structure
 4090 referenced by the *aiocbp* argument. The error status for an asynchronous I/O operation is the
 4091 SIO *errno* value that would be set by the corresponding *read()*, *write()*, *fdatasync()*, or *fsync()*
 4092 operation. If the operation has not yet completed, then the error status shall be equal to
 4093 [EINPROGRESS].

4094 **RETURN VALUE**

4095 If the asynchronous I/O operation has completed successfully, then 0 shall be returned. If the
 4096 asynchronous operation has completed unsuccessfully, then the error status, as described for
 4097 SIO *read()*, *write()*, *fdatasync()*, and *fsync()*, shall be returned. If the asynchronous I/O operation has
 4098 not yet completed, then [EINPROGRESS] shall be returned.

4099 **ERRORS**4100 The *aio_error()* function may fail if:

4101 [EINVAL] The *aiocbp* argument does not refer to an asynchronous operation whose
 4102 return status has not yet been retrieved.

4103 **EXAMPLES**

4104 None.

4105 **APPLICATION USAGE**

4106 The *aio_error()* function is part of the Asynchronous Input and Output option and need not be
 4107 available on all implementations.

4108 **RATIONALE**

4109 None.

4110 **FUTURE DIRECTIONS**

4111 None.

4112 **SEE ALSO**

4113 *aio_cancel()*, *aio_fsync()*, *aio_read()*, *aio_return()*, *aio_write()*, *close()*, *exec*, *exit()*, *fork()*, *lio_listio()*,
 4114 *lseek()*, *read()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**aio.h**>

4115 **CHANGE HISTORY**

4116 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4117 **Issue 6**

4118 The [ENOSYS] error condition has been removed as stubs need not be provided if an
 4119 implementation does not support the Asynchronous Input and Output option.

4120 The APPLICATION USAGE section is added.

4121 NAME

4122 aio_fsync — asynchronous file synchronization (**REALTIME**)

4123 SYNOPSIS

4124 AIO #include <aio.h>

4125 int aio_fsync(int op, struct aiocb *aiocbp);

4126

4127 DESCRIPTION

4128 The *aio_fsync()* function asynchronously forces all I/O operations associated with the file
 4129 indicated by the file descriptor *aio_fildes* member of the **aiocb** structure referenced by the *aiocbp*
 4130 argument and queued at the time of the call to *aio_fsync()* to the synchronized I/O completion
 4131 state. The function call shall return when the synchronization request has been initiated or
 4132 queued to the file or device (even when the data cannot be synchronized immediately).

4133 If *op* is O_DSYNC, all currently queued I/O operations shall be completed as if by a call to
 4134 *fdatasync()*; that is, as defined for synchronized I/O data integrity completion. If *op* is O_SYNC,
 4135 all currently queued I/O operations shall be completed as if by a call to *fsync()*; that is, as
 4136 defined for synchronized I/O file integrity completion. If the *aio_fsync()* function fails, or if the
 4137 operation queued by *aio_fsync()* fails, then, as for *fsync()* and *fdatasync()*, outstanding I/O
 4138 operations are not guaranteed to have been completed.

4139 If *aio_fsync()* succeeds, then it is only the I/O that was queued at the time of the call to
 4140 *aio_fsync()* that is guaranteed to be forced to the relevant completion state. The completion of
 4141 subsequent I/O on the file descriptor is not guaranteed to be completed in a synchronized
 4142 fashion.

4143 The *aiocbp* argument refers to an asynchronous I/O control block. The *aiocbp* value may be used
 4144 as an argument to *aio_error()* and *aio_return()* in order to determine the error status and return
 4145 status, respectively, of the asynchronous operation while it is proceeding. When the request is
 4146 queued, the error status for the operation is [EINPROGRESS]. When all data has been
 4147 successfully transferred, the error status shall be reset to reflect the success or failure of the
 4148 operation. If the operation does not complete successfully, the error status for the operation shall
 4149 be set to indicate the error. The *aio_sigevent* member determines the asynchronous notification to
 4150 occur as specified in Section 2.4.1 (on page 528) when all operations have achieved synchronized
 4151 I/O completion. All other members of the structure referenced by *aiocbp* are ignored. If the
 4152 control block referenced by *aiocbp* becomes an illegal address prior to asynchronous I/O
 4153 completion, then the behavior is undefined.

4154 If the *aio_fsync()* function fails or the *aiocbp* indicates an error condition, data is not guaranteed
 4155 to have been successfully transferred.

4156 RETURN VALUE

4157 The *aio_fsync()* function shall return the value 0 to the calling process if the I/O operation is
 4158 successfully queued; otherwise, the function shall return the value -1 and set *errno* to indicate
 4159 the error.

4160 ERRORS

4161 The *aio_fsync()* function shall fail if:

4162 [EAGAIN] The requested asynchronous operation was not queued due to temporary
 4163 resource limitations.

4164 [EBADF] The *aio_fildes* member of the **aiocb** structure referenced by the *aiocbp* argument
 4165 is not a valid file descriptor open for writing.

4166 [EINVAL] This implementation does not support synchronized I/O for this file. |
4167 [EINVAL] A value of *op* other than O_DSYNC or O_SYNC was specified.
4168 In the event that any of the queued I/O operations fail, *aio_fsync()* shall return the error
4169 condition defined for *read()* and *write()*. The error is returned in the error status for the
4170 asynchronous *fsync()* operation, which can be retrieved using *aio_error()*.

4171 EXAMPLES

4172 None.

4173 APPLICATION USAGE

4174 The *aio_fsync()* function is part of the Asynchronous Input and Output option and need not be |
4175 available on all implementations.

4176 RATIONALE

4177 None.

4178 FUTURE DIRECTIONS

4179 None.

4180 SEE ALSO

4181 *fcntl()*, *fdatasync()*, *fsync()*, *open()*, *read()*, *write()*, the Base Definitions volume of |
4182 IEEE Std. 1003.1-200x, <**aio.h**>

4183 CHANGE HISTORY

4184 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4185 Issue 6

4186 The [ENOSYS] error condition has been removed as stubs need not be provided if an |
4187 implementation does not support the Asynchronous Input and Output option.

4188 The APPLICATION USAGE section is added.

4189 **NAME**4190 aio_read — asynchronous read from a file (**REALTIME**)4191 **SYNOPSIS**

4192 AIO #include <aio.h>

4193 int aio_read(struct aiocb *aiocbp);

4194

4195 **DESCRIPTION**

4196 The *aio_read()* function allows the calling process to read *aiocbp->aio_nbytes* from the file
 4197 associated with *aiocbp->aio_fildes* into the buffer pointed to by *aiocbp->aio_buf*. The function call
 4198 shall return when the read request has been initiated or queued to the file or device (even when
 4199 the data cannot be delivered immediately).

4200 PIO If prioritized I/O is supported for this file, then the asynchronous operation is submitted at a
 4201 priority equal to the scheduling priority of the process minus *aiocbp->aio_reqprio*.

4202 The *aiocbp* value may be used as an argument to *aio_error()* and *aio_return()* in order to
 4203 determine the error status and return status, respectively, of the asynchronous operation while it
 4204 is proceeding. If an error condition is encountered during queuing, the function call shall return
 4205 without having initiated or queued the request. The requested operation takes place at the
 4206 absolute position in the file as given by *aio_offset*, as if *lseek()* were called immediately prior to
 4207 the operation with an *offset* equal to *aio_offset* and a *whence* equal to {SEEK_SET}. After a
 4208 successful call to enqueue an asynchronous I/O operation, the value of the file offset for the file
 4209 is unspecified.

4210 The *aiocbp->aio_lio_opcode* field shall be ignored by *aio_read()*.

4211 The *aiocbp* argument points to an **aiocb** structure. If the buffer pointed to by *aiocbp->aio_buf* or
 4212 the control block pointed to by *aiocbp* becomes an illegal address prior to asynchronous I/O
 4213 completion, then the behavior is undefined.

4214 Simultaneous asynchronous operations using the same *aiocbp* produce undefined results.

4215 SIO If synchronized I/O is enabled on the file associated with *aiocbp->aio_fildes*, the behavior of this
 4216 function shall be according to the definitions of synchronized I/O data integrity completion and
 4217 synchronized I/O file integrity completion.

4218 For any system action that changes the process memory space while an asynchronous I/O is
 4219 outstanding to the address range being changed, the result of that action is undefined.

4220 For regular files, no data transfer shall occur past the offset maximum established in the open
 4221 file description associated with *aiocbp->aio_fildes*.

4222 **RETURN VALUE**

4223 The *aio_read()* function shall return the value zero to the calling process if the I/O operation is
 4224 successfully queued; otherwise, the function shall return the value -1 and set *errno* to indicate
 4225 the error.

4226 **ERRORS**

4227 The *aio_read()* function shall fail if:

4228 [EAGAIN] The requested asynchronous I/O operation was not queued due to system
 4229 resource limitations.

4230 Each of the following conditions may be detected synchronously at the time of the call to
 4231 *aio_read()*, or asynchronously. If any of the conditions below are detected synchronously, the
 4232 *aio_read()* function shall return -1 and set *errno* to the corresponding value. If any of the
 4233 conditions below are detected asynchronously, the return status of the asynchronous operation

- 4234 is set to `-1`, and the error status of the asynchronous operation is set to the corresponding value.
- 4235 [EBADF] The `aiocbp->aio_fildes` argument is not a valid file descriptor open for reading.
- 4236 [EINVAL] The file offset value implied by `aiocbp->aio_offset` would be invalid, `aiocbp->aio_reqprio` is not a valid value, or `aiocbp->aio_nbytes` is an invalid value.
- 4237
- 4238 In the case that the `aio_read()` successfully queues the I/O operation but the operation is subsequently canceled or encounters an error, the return status of the asynchronous operation is one of the values normally returned by the `read()` function call. In addition, the error status of the asynchronous operation is set to one of the error statuses normally set by the `read()` function call, or one of the following values:
- 4239
- 4240
- 4241
- 4242
- 4243 [EBADF] The `aiocbp->aio_fildes` argument is not a valid file descriptor open for reading.
- 4244 [ECANCELED] The requested I/O was canceled before the I/O completed due to an explicit `aio_cancel()` request.
- 4245
- 4246 [EINVAL] The file offset value implied by `aiocbp->aio_offset` would be invalid.
- 4247 The following condition may be detected synchronously or asynchronously:
- 4248 [EOVERFLOW] The file is a regular file, `aiocbp->aio_nbytes` is greater than 0, and the starting offset in `aiocbp->aio_offset` is before the end-of-file and is at or beyond the offset maximum in the open file description associated with `aiocbp->aio_fildes`.
- 4249
- 4250
- 4251 **EXAMPLES**
- 4252 None.
- 4253 **APPLICATION USAGE**
- 4254 The `aio_read()` function is part of the Asynchronous Input and Output option and need not be available on all implementations.
- 4255
- 4256 **RATIONALE**
- 4257 None.
- 4258 **FUTURE DIRECTIONS**
- 4259 None.
- 4260 **SEE ALSO**
- 4261 `aio_cancel()`, `aio_error()`, `lio_listio()`, `aio_return()`, `aio_write()`, `close()`, `exec`, `exit()`, `fork()`, `lseek()`, `read()`, the Base Definitions volume of IEEE Std. 1003.1-200x, <`aio.h`>
- 4262
- 4263 **CHANGE HISTORY**
- 4264 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.
- 4265 Large File Summit extensions are added.
- 4266 **Issue 6**
- 4267 The [ENOSYS] error condition has been removed as stubs need not be provided if an implementation does not support the Asynchronous Input and Output option.
- 4268
- 4269 The APPLICATION USAGE section is added.
- 4270 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:
- 4271
- 4272 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file description. This change is to support large files.
 - 4273
 - 4274 • In the ERRORS section, the [EOVERFLOW] condition is added. This change is to support large files.
 - 4275

4276 **NAME**4277 aio_return — retrieve return status of an asynchronous I/O operation (**REALTIME**)4278 **SYNOPSIS**4279 AIO `#include <aio.h>`4280 `ssize_t aio_return(struct aiocb *aiocbp);`

4281

4282 **DESCRIPTION**

4283 The `aio_return()` function shall return the return status associated with the **aiocb** structure
 4284 referenced by the `aiocbp` argument. The return status for an asynchronous I/O operation is the
 4285 value that would be returned by the corresponding `read()`, `write()`, or `fsync()` function call. If the
 4286 error status for the operation is equal to [EINPROGRESS], then the return status for the
 4287 operation is undefined. The `aio_return()` function may be called exactly once to retrieve the
 4288 return status of a given asynchronous operation; thereafter, if the same **aiocb** structure is used in
 4289 a call to `aio_return()` or `aio_error()`, an error may be returned. When the **aiocb** structure referred
 4290 to by `aiocbp` is used to submit another asynchronous operation, then `aio_return()` may be
 4291 successfully used to retrieve the return status of that operation.

4292 **RETURN VALUE**

4293 If the asynchronous I/O operation has completed, then the return status, as described for `read()`,
 4294 `write()`, and `fsync()`, shall be returned. If the asynchronous I/O operation has not yet completed,
 4295 the results of `aio_return()` are undefined.

4296 **ERRORS**4297 The `aio_return()` function may fail if:

4298 [EINVAL] The `aiocbp` argument does not refer to an asynchronous operation whose
 4299 return status has not yet been retrieved.

4300 **EXAMPLES**

4301 None.

4302 **APPLICATION USAGE**

4303 The `aio_return()` function is part of the Asynchronous Input and Output option and need not be
 4304 available on all implementations.

4305 **RATIONALE**

4306 None.

4307 **FUTURE DIRECTIONS**

4308 None.

4309 **SEE ALSO**

4310 `aio_cancel()`, `aio_error()`, `aio_fsync()`, `aio_read()`, `aio_write()`, `close()`, `exec`, `exit()`, `fork()`, `lio_listio()`,
 4311 `lseek()`, `read()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<aio.h>`

4312 **CHANGE HISTORY**

4313 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4314 **Issue 6**

4315 The [ENOSYS] error condition has been removed as stubs need not be provided if an
 4316 implementation does not support the Asynchronous Input and Output option.

4317 The APPLICATION USAGE section is added.

4318 The [EINVAL] error condition is updated as a “may fail”. This is for consistency with the
 4319 DESCRIPTION.

4320 NAME

4321 aio_suspend — wait for an asynchronous I/O request (**REALTIME**)

4322 SYNOPSIS

4323 AIO #include <aio.h>

```
4324 int aio_suspend(const struct aiocb * const list[], int nent,
4325               const struct timespec *timeout);
4326
```

4327 DESCRIPTION

4328 The *aio_suspend()* function shall suspend the calling thread until at least one of the asynchronous
 4329 I/O operations referenced by the *list* argument has completed, until a signal interrupts the
 4330 function, or, if *timeout* is not NULL, until the time interval specified by *timeout* has passed. If any
 4331 of the **aiocb** structures in the list correspond to completed asynchronous I/O operations (that is,
 4332 the error status for the operation is not equal to [EINPROGRESS]) at the time of the call, the
 4333 function shall return without suspending the calling thread. The *list* argument is an array of
 4334 pointers to asynchronous I/O control blocks. The *nent* argument indicates the number of
 4335 elements in the array. Each **aiocb** structure pointed to has been used in initiating an
 4336 asynchronous I/O request via *aio_read()*, *aio_write()*, or *lio_listio()*. This array may contain
 4337 NULL pointers, which are ignored. If this array contains pointers that refer to **aiocb** structures
 4338 that have not been used in submitting asynchronous I/O, the effect is undefined.

4339 If the time interval indicated in the **timespec** structure pointed to by *timeout* passes before any of
 4340 the I/O operations referenced by *list* are completed, then *aio_suspend()* shall return with an
 4341 error. If the Monotonic Clock option is supported, the clock that shall be used to measure this
 4342 time interval shall be the CLOCK_MONOTONIC clock.

4343 RETURN VALUE

4344 If the *aio_suspend()* function returns after one or more asynchronous I/O operations have
 4345 completed, the function shall return zero. Otherwise, the function shall return a value of -1 and
 4346 set *errno* to indicate the error.

4347 The application may determine which asynchronous I/O completed by scanning the associated
 4348 error and return status using *aio_error()* and *aio_return()*, respectively.

4349 ERRORS

4350 The *aio_suspend()* function shall fail if:

4351 [EAGAIN] No asynchronous I/O indicated in the list referenced by *list* completed in the
 4352 time interval indicated by *timeout*.

4353 [EINTR] A signal interrupted the *aio_suspend()* function. Note that, since each
 4354 asynchronous I/O operation may possibly provoke a signal when it
 4355 completes, this error return may be caused by the completion of one (or more)
 4356 of the very I/O operations being awaited.

4357 EXAMPLES

4358 None.

4359 APPLICATION USAGE

4360 The *aio_suspend()* function is part of the Asynchronous Input and Output option and need not
 4361 be available on all implementations.

4362 RATIONALE

4363 None.

4364 **FUTURE DIRECTIONS**

4365 None.

4366 **SEE ALSO**4367 *aio_read()*, *aio_write()*, *lio_listio()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**aio.h**>4368 **CHANGE HISTORY**

4369 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4370 **Issue 6**4371 The [ENOSYS] error condition has been removed as stubs need not be provided if an
4372 implementation does not support the Asynchronous Input and Output option.

4373 The APPLICATION USAGE section is added.

4374 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that the
4375 CLOCK_MONOTONIC clock, if supported, is used.

4376 NAME

4377 aio_write — asynchronous write to a file (**REALTIME**)

4378 SYNOPSIS

4379 AIO #include <aio.h>

4380 int aio_write(struct aiocb *aiocbp);

4381

4382 DESCRIPTION

4383 The *aio_write()* function allows the calling process to write *aiocbp->aio_nbytes* to the file
 4384 associated with *aiocbp->aio_fildes* from the buffer pointed to by *aiocbp->aio_buf*. The function call
 4385 shall return when the write request has been initiated or, at a minimum, queued to the file or
 4386 device.

4387 PIO If prioritized I/O is supported for this file, then the asynchronous operation shall be submitted
 4388 at a priority equal to the scheduling priority of the process minus *aiocbp->aio_reqprio*.

4389 The *aiocbp* may be used as an argument to *aio_error()* and *aio_return()* in order to determine the
 4390 error status and return status, respectively, of the asynchronous operation while it is proceeding.

4391 The *aiocbp* argument points to an **aiocb** structure. If the buffer pointed to by *aiocbp->aio_buf* or
 4392 the control block pointed to by *aiocbp* becomes an illegal address prior to asynchronous I/O
 4393 completion, then the behavior is undefined.

4394 If **O_APPEND** is not set for the file descriptor *aio_fildes*, then the requested operation takes place
 4395 at the absolute position in the file as given by *aio_offset*, as if *lseek()* were called immediately
 4396 prior to the operation with an *offset* equal to *aio_offset* and a *whence* equal to **{SEEK_SET}**. If
 4397 **O_APPEND** is set for the file descriptor, write operations append to the file in the same order as
 4398 the calls were made. After a successful call to enqueue an asynchronous I/O operation, the value
 4399 of the file offset for the file is unspecified.

4400 The *aiocbp->aio_lio_opcode* field shall be ignored by *aio_write()*.

4401 Simultaneous asynchronous operations using the same *aiocbp* produce undefined results.

4402 SIO If synchronized I/O is enabled on the file associated with *aiocbp->aio_fildes*, the behavior of this
 4403 function shall be according to the definitions of synchronized I/O data integrity completion, and
 4404 synchronized I/O file integrity completion.

4405 For any system action that changes the process memory space while an asynchronous I/O is
 4406 outstanding to the address range being changed, the result of that action is undefined.

4407 For regular files, no data transfer shall occur past the offset maximum established in the open
 4408 file description associated with *aiocbp->aio_fildes*.

4409 RETURN VALUE

4410 The *aio_write()* function shall return the value zero to the calling process if the I/O operation is
 4411 successfully queued; otherwise, the function shall return the value **-1** and set *errno* to indicate
 4412 the error.

4413 ERRORS

4414 The *aio_write()* function shall fail if:

4415 [EAGAIN] The requested asynchronous I/O operation was not queued due to system
 4416 resource limitations.

4417 Each of the following conditions may be detected synchronously at the time of the call to
 4418 *aio_write()*, or asynchronously. If any of the conditions below are detected synchronously, the
 4419 *aio_write()* function shall return **-1** and set *errno* to the corresponding value. If any of the

4420 conditions below are detected asynchronously, the return status of the asynchronous operation
 4421 shall be set to -1, and the error status of the asynchronous operation is set to the corresponding
 4422 value.

4423 [EBADF] The *aiocbp->aio_fildes* argument is not a valid file descriptor open for writing.

4424 [EINVAL] The file offset value implied by *aiocbp->aio_offset* would be invalid, *aiocbp-*
 4425 *>aio_reqprio* is not a valid value, or *aiocbp->aio_nbytes* is an invalid value.

4426 In the case that the *aio_write()* successfully queues the I/O operation, the return status of the
 4427 asynchronous operation shall be one of the values normally returned by the *write()* function call.
 4428 If the operation is successfully queued but is subsequently canceled or encounters an error, the
 4429 error status for the asynchronous operation contains one of the values normally set by the
 4430 *write()* function call, or one of the following:

4431 [EBADF] The *aiocbp->aio_fildes* argument is not a valid file descriptor open for writing.

4432 [EINVAL] The file offset value implied by *aiocbp->aio_offset* would be invalid.

4433 [ECANCELED] The requested I/O was canceled before the I/O completed due to an explicit
 4434 *aio_cancel()* request.

4435 The following condition may be detected synchronously or asynchronously:

4436 [EFBIG] The file is a regular file, *aiocbp->aio_nbytes* is greater than 0, and the starting
 4437 offset in *aiocbp->aio_offset* is at or beyond the offset maximum in the open file
 4438 description associated with *aiocbp->aio_fildes*.

4439 EXAMPLES

4440 None.

4441 APPLICATION USAGE

4442 The *aio_write()* function is part of the Asynchronous Input and Output option and need not be
 4443 available on all implementations.

4444 RATIONALE

4445 None.

4446 FUTURE DIRECTIONS

4447 None.

4448 SEE ALSO

4449 *aio_cancel()*, *aio_error()*, *aio_read()*, *aio_return()*, *close()*, *exec*, *exit()*, *fork()*, *lio_listio()*, *lseek()*,
 4450 *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**aio.h**>

4451 CHANGE HISTORY

4452 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

4453 Large File Summit extensions are added.

4454 Issue 6

4455 The [ENOSYS] error condition has been removed as stubs need not be provided if an
 4456 implementation does not support the Asynchronous Input and Output option.

4457 The APPLICATION USAGE section is added.

4458 The following new requirements on POSIX implementations derive from alignment with the
 4459 Single UNIX Specification:

- 4460 • In the DESCRIPTION, text is added to indicate that for regular files no data transfer occurs
 4461 past the offset maximum established in the open file description associated with *aiocbp-*
 4462 *>aio_fildes*.

4463

- The [EFBIG] error is added as part of the large file support extensions.

4464 **NAME**

4465 alarm — schedule an alarm signal

4466 **SYNOPSIS**

4467 #include <unistd.h>

4468 unsigned alarm(unsigned *seconds*);4469 **DESCRIPTION**

4470 The *alarm()* function shall cause the system to generate a SIGALRM signal for the process after
4471 the number of realtime seconds specified by *seconds* have elapsed. Processor scheduling delays
4472 may prevent the process from handling the signal as soon as it is generated.

4473 If *seconds* is 0, a pending alarm request, if any, is canceled.

4474 Alarm requests are not stacked; only one SIGALRM generation can be scheduled in this manner.
4475 If the SIGALRM signal has not yet been generated, the call shall result in rescheduling the time
4476 at which the SIGALRM signal is generated.

4477 Interactions between *alarm()* and any of *setitimer()*, *ualarm()*, or *usleep()* are unspecified.

4478 **RETURN VALUE**

4479 If there is a previous *alarm()* request with time remaining, *alarm()* shall return a non-zero value
4480 that is the number of seconds until the previous request would have generated a SIGALRM
4481 signal. Otherwise, *alarm()* shall return 0.

4482 **ERRORS**

4483 The *alarm()* function is always successful, and no return value is reserved to indicate an error.

4484 **EXAMPLES**

4485 None.

4486 **APPLICATION USAGE**

4487 The *fork()* function clears pending alarms in the child process. A new process image created by
4488 one of the *exec* functions inherits the time left to an alarm signal in the old process' image.

4489 Application writers should note that the type of the argument *seconds* and the return value of
4490 *alarm()* is **unsigned**. That means that a Strictly Conforming POSIX System Interfaces
4491 Application cannot pass a value greater than the minimum guaranteed value for {UINT_MAX},
4492 which the ISO C standard sets as 65 535, and any application passing a larger value is restricting
4493 its portability. A different type was considered, but historical implementations, including those
4494 with a 16-bit **int** type, consistently use either **unsigned** or **int**.

4495 Application writers should be aware of possible interactions when the same process uses both
4496 the *alarm()* and *sleep()* functions.

4497 **RATIONALE**

4498 Many historical implementations (including Version 7 and System V) allow an alarm to occur up
4499 to a second early. Other implementations allow alarms up to half a second or one clock tick
4500 early or do not allow them to occur early at all. The latter is considered most appropriate, since it
4501 gives the most predictable behavior, especially since the signal can always be delayed for an
4502 indefinite amount of time due to scheduling. Applications can thus choose the *seconds* argument
4503 as the minimum amount of time they wish to have elapse before the signal.

4504 The term *realtime* here and elsewhere (*sleep()*, *times()*) is intended to mean “wall clock” time as
4505 common English usage, and has nothing to do with “realtime operating systems”. It is in
4506 contrast to *virtual time*, which could be misinterpreted if just *time* were used.

4507 In some implementations, including 4.3 BSD, very large values of the *seconds* argument are
4508 silently rounded down to an implementation-defined maximum value. This maximum is large

4509 enough (on the order of several months) that the effect is not noticeable.

4510 There were two possible choices for alarm generation in multi-threaded applications: generation
4511 for the calling thread or generation for the process. The first option would not have been
4512 particularly useful since the alarm state is maintained on a per-process basis and the alarm that
4513 is established by the last invocation of *alarm()* is the only one that would be active.

4514 Furthermore, allowing generation of an asynchronous signal for a thread would have introduced
4515 an exception to the overall signal model. This requires a compelling reason in order to be
4516 justified.

4517 **FUTURE DIRECTIONS**

4518 None.

4519 **SEE ALSO**

4520 *alarm()*, *exec*, *fork()*, *getitimer()*, *pause()*, *sigaction()*, *sleep()*, *ualarm()*, *usleep()*, the Base
4521 Definitions volume of IEEE Std. 1003.1-200x, <**signal.h**>, <**unistd.h**>

4522 **CHANGE HISTORY**

4523 First released in Issue 1. Derived from Issue 1 of the SVID.

4524 **Issue 4**

4525 The <**unistd.h**> header is included in the SYNOPSIS section.

4526 **Issue 4, Version 2**

4527 The DESCRIPTION is updated to indicate that interactions with the *setitimer()*, *ualarm()*, and
4528 *usleep()* functions are unspecified.

4529 **Issue 6**

4530 The following new requirements on POSIX implementations derive from alignment with the
4531 Single UNIX Specification:

- 4532 • The DESCRIPTION is updated to indicate that interactions with the *setitimer()*, *ualarm()*, and
4533 *usleep()* functions are unspecified.

4534 **NAME**

4535 asctime, asctime_r — convert date and time to a string

4536 **SYNOPSIS**

4537 #include <time.h>

4538 char *asctime(const struct tm *timeptr);

4539 TSF char *asctime_r(const struct tm *restrict tm, char *restrict buf);

4540

4541 **DESCRIPTION**

4542 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 4543 conflict between the requirements described here and the ISO C standard is unintentional. This
 4544 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

4545 The *asctime()* function shall convert the broken-down time in the structure pointed to by *timeptr*
 4546 into a string in the form:

4547 Sun Sep 16 01:03:52 1973\n\0

4548 using the equivalent of the following algorithm:

```

4549 char *asctime(const struct tm *timeptr)
4550 {
4551     static char wday_name[7][3] = {
4552         "Sun", "Mon", "Tue", "Wed", "Thu", "Fri", "Sat"
4553     };
4554     static char mon_name[12][3] = {
4555         "Jan", "Feb", "Mar", "Apr", "May", "Jun",
4556         "Jul", "Aug", "Sep", "Oct", "Nov", "Dec"
4557     };
4558     static char result[26];

4559     sprintf(result, "%.3s %.3s%3d %.2d:%.2d:%.2d %d\n",
4560             wday_name[timeptr->tm_wday],
4561             mon_name[timeptr->tm_mon],
4562             timeptr->tm_mday, timeptr->tm_hour,
4563             timeptr->tm_min, timeptr->tm_sec,
4564             1900 + timeptr->tm_year);
4565     return result;
4566 }
```

4567 The **tm** structure is defined in the **<time.h>** header.

4568 CX The *asctime()*, *ctime()*, *gmtime()*, and *localtime()* functions shall return values in one of two static
 4569 objects: a broken-down time structure and an array of type **char**. Execution of any of the
 4570 functions may overwrite the information returned in either of these objects by any of the other
 4571 functions.

4572 The *asctime()* function need not be reentrant. A function that is not required to be reentrant is not
 4573 required to be thread-safe.

4574 TSF The *asctime_r()* function shall convert the broken-down time in the structure pointed to by *tm*
 4575 into a string (of the same form as that returned by *asctime()*) that is placed in the user-supplied
 4576 buffer pointed to by *buf* (which contains at least 26 bytes) and then return *buf*.

4577 **RETURN VALUE**

4578 Upon successful completion, *asctime()* shall return a pointer to the string.

4579 TSF Upon successful completion, *asctime_r()* shall return a pointer to a character string containing
4580 the date and time. This string is pointed to by the argument *buf*. If the function is unsuccessful,
4581 it shall return NULL.

4582 **ERRORS**

4583 No errors are defined.

4584 **EXAMPLES**

4585 None.

4586 **APPLICATION USAGE**

4587 Values for the broken-down time structure can be obtained by calling *gmtime()* or *localtime()*.
4588 This function is included for compatibility with older implementations, and does not support
4589 localized date and time formats. Applications should use *strptime()* to achieve maximum
4590 portability.

4591 The *asctime_r()* function is thread-safe and shall return values in a user-supplied buffer instead
4592 of possibly using a static data area that may be overwritten by each call.

4593 **RATIONALE**

4594 None.

4595 **FUTURE DIRECTIONS**

4596 None.

4597 **SEE ALSO**

4598 *clock()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strptime()*, *strptime()*, *time()*, *utime()*,
4599 the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

4600 **CHANGE HISTORY**

4601 First released in Issue 1. Derived from Issue 1 of the SVID.

4602 **Issue 4**

4603 The location of the **tm** structure is now defined.

4604 The APPLICATION USAGE section is expanded to describe the time-handling functions
4605 generally and to refer users to *strptime()*, which is a locale-dependent time-handling function.

4606 The following change is incorporated for alignment with the ISO C standard:

- 4607 • The type of argument *timeptr* is changed from **struct tm*** to **const struct tm***.

4608 **Issue 5**

4609 Normative text previously in the APPLICATION USAGE section is moved to the
4610 DESCRIPTION.

4611 The *asctime_r()* function is included for alignment with the POSIX Threads Extension.

4612 A note indicating that the *asctime()* function need not be reentrant is added to the
4613 DESCRIPTION.

4614 **Issue 6**

4615 The *asctime_r()* function is marked as part of the Thread-Safe Functions option.

4616 Extensions beyond the ISO C standard are now marked.

4617 The APPLICATION USAGE section is updated to include a note on the thread-safe function and
4618 its avoidance of possibly using a static data area.

4619 The DESCRIPTION of *asctime_r()* is updated to describe the format of the string returned. |
4620 The **restrict** keyword is added to the *asctime_r()* prototype for alignment with the |
4621 ISO/IEC 9899:1999 standard. |

4622 **NAME**

4623 asin, asinf, asinl — arc sine function

4624 **SYNOPSIS**

4625 #include <math.h>

4626 double asin(double x);

4627 float asinf(float x);

4628 long double asinl(long double x);

4629 **DESCRIPTION**

4630 cx The functionality described on this reference page is aligned with the ISO C standard. Any
 4631 conflict between the requirements described here and the ISO C standard is unintentional. This
 4632 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

4633 The *asin()*, *asinf()*, and *asinl()* functions shall compute the principal value of the arc sine of *x*.
 4634 The value of *x* should be in the range $[-1,1]$.

4635 An application wishing to check for error situations should set *errno* to 0, then call *asin()*. If *errno*
 4636 is non-zero on return, or the return value is NaN, an error has occurred.

4637 **RETURN VALUE**

4638 Upon successful completion, the *asin()*, *asinf()*, and *asinl()* functions shall return the arc sine of
 4639 xsi *x*, in the range $[-\pi/2, \pi/2]$ radians. If the value of *x* is not in the range $[-1,1]$, and is not $\pm\text{Inf}$ or
 4640 NaN, either 0.0 or NaN shall be returned and *errno* shall be set to [EDOM].

4641 xsi If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

4642 If *x* is $\pm\text{Inf}$, either 0.0 shall be returned and *errno* set to [EDOM], or NaN shall be returned and
 4643 *errno* may be set to [EDOM].

4644 If the result underflows, 0.0 shall be returned and *errno* may be set to [ERANGE].

4645 **ERRORS**

4646 The *asin()*, *asinf()*, and *asinl()* functions shall fail if:

4647 xsi [EDOM] The value *x* is not $\pm\text{Inf}$ or NaN and is not in the range $[-1,1]$.

4648 The *asin()*, *asinf()*, and *asinl()* functions may fail if:

4649 xsi [EDOM] The value of *x* is $\pm\text{Inf}$ or NaN.

4650 [ERANGE] The result underflows

4651 xsi No other errors shall occur.

4652 **EXAMPLES**

4653 None.

4654 **APPLICATION USAGE**

4655 None.

4656 **RATIONALE**

4657 None.

4658 **FUTURE DIRECTIONS**

4659 None.

4660 **SEE ALSO**

4661 *isnan()*, *sin()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

4662 **CHANGE HISTORY**

4663 First released in Issue 1. Derived from Issue 1 of the SVID.

4664 **Issue 4**

4665 References to *matherr()* are removed.

4666 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the
4667 ISO C standard and to rationalize error handling in the mathematics functions.

4668 The return value specified for [EDOM] is marked as an extension.

4669 **Issue 5**

4670 The DESCRIPTION is updated to indicate how an application should check for an error. This
4671 text was previously published in the APPLICATION USAGE section.

4672 **Issue 6**

4673 The *asinf()* and *asinl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

4674 **NAME**

4675 asinh — hyperbolic arc sine

4676 **SYNOPSIS**

4677 xSI #include <math.h>

4678 double asinh(double x);

4679

4680 **DESCRIPTION**

4681 Refer to *acosh()*.

4682 **NAME**

4683 assert — insert program diagnostics

4684 **SYNOPSIS**

4685 #include <assert.h>

4686 void assert(*scalar expression*);4687 **DESCRIPTION**

4688 **cx** The functionality described on this reference page is aligned with the ISO C standard. Any
4689 conflict between the requirements described here and the ISO C standard is unintentional. This
4690 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

4691 The *assert()* macro inserts diagnostics into programs; it expands to a **void** expression. When it is
4692 executed, if *expression* (which shall have a **scalar** type) is false (that is, compares equal to 0),
4693 *assert()* shall write information about the particular call that failed on *stderr* and shall call *abort()*.

4694 The information written about the call that failed shall include the text of the argument, the
4695 name of the source file, the source file line number, and the name of the enclosing function, the
4696 latter are, respectively, the values of the preprocessing macros `__FILE__` and `__LINE__` and of
4697 the identifier `__func__`.

4698 Forcing a definition of the name `NDEBUG`, either from the compiler command line or with the
4699 preprocessor control statement `#define NDEBUG` ahead of the `#include <assert.h>` statement,
4700 shall stop assertions from being compiled into the program.

4701 **RETURN VALUE**4702 The *assert()* macro shall return no value.4703 **ERRORS**

4704 No errors are defined.

4705 **EXAMPLES**

4706 None.

4707 **APPLICATION USAGE**

4708 None.

4709 **RATIONALE**

4710 None.

4711 **FUTURE DIRECTIONS**

4712 None.

4713 **SEE ALSO**4714 *abort()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<assert.h>`, *stderr*4715 **CHANGE HISTORY**

4716 First released in Issue 1. Derived from Issue 1 of the SVID.

4717 **Issue 4**

4718 The APPLICATION USAGE section is merged into the DESCRIPTION.

4719 **Issue 6**4720 The prototype for the *expression* argument to *assert()* is changed from **int** to **scalar** for alignment
4721 with the ISO/IEC 9899:1999 standard.4722 The DESCRIPTION of *assert()* is updated for alignment with the ISO/IEC 9899:1999 standard.

4723 **NAME**

4724 atan, atanf, atanl — arc tangent function

4725 **SYNOPSIS**

4726 #include <math.h>

4727 double atan(double x);

4728 float atanf(float x);

4729 long double atanl(long double x);

4730 **DESCRIPTION**

4731 cx The functionality described on this reference page is aligned with the ISO C standard. Any
4732 conflict between the requirements described here and the ISO C standard is unintentional. This
4733 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

4734 The *atan()*, *atanf()*, and *atanl()* functions shall compute the principal value of the arc tangent of
4735 *x*.

4736 An application wishing to check for error situations should set *errno* to 0 before calling *atan()*. If
4737 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

4738 **RETURN VALUE**

4739 Upon successful completion, the *atan()*, *atanf()*, and *atanl()* functions shall return the arc
4740 tangent of *x* in the range $[-\pi/2, \pi/2]$ radians.

4741 xSI If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

4742 If the result underflows, 0.0 shall be returned and *errno* may be set to [ERANGE].

4743 **ERRORS**

4744 The *atan()*, *atanf()*, and *atanl()* functions may fail if:

4745 xSI [EDOM] The value of *x* is NaN.

4746 [ERANGE] The result underflows

4747 xSI No other errors shall occur.

4748 **EXAMPLES**

4749 None.

4750 **APPLICATION USAGE**

4751 None.

4752 **RATIONALE**

4753 None.

4754 **FUTURE DIRECTIONS**

4755 None.

4756 **SEE ALSO**

4757 *atan2()*, *isnan()*, *tan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

4758 **CHANGE HISTORY**

4759 First released in Issue 1. Derived from Issue 1 of the SVID.

4760 **Issue 4**

4761 References to *matherr()* are removed.

4762 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the
4763 ISO C standard and to rationalize error handling in the mathematics functions.

- 4764 The return value specified for [EDOM] is marked as an extension.
- 4765 **Issue 5**
- 4766 The DESCRIPTION is updated to indicate how an application should check for an error. This
- 4767 text was previously published in the APPLICATION USAGE section.
- 4768 **Issue 6**
- 4769 The *atanf()* and *atanl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

4770 **NAME**4771 `atan2` — arc tangent function4772 **SYNOPSIS**4773 `#include <math.h>`4774 `double atan2(double y, double x);`4775 `float atan2f(float y, float x);`4776 `long double atan2l(long double y, long double x);`4777 **DESCRIPTION**

4778 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
 4779 conflict between the requirements described here and the ISO C standard is unintentional. This
 4780 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

4781 The `atan2()`, `atan2f()`, and `atan2l()` functions shall compute the principal value of the arc tangent
 4782 of y/x , using the signs of both arguments to determine the quadrant of the return value.

4783 An application wishing to check for error situations should set `errno` to 0 before calling `atan2()`.
 4784 If `errno` is non-zero on return, or the return value is NaN, an error has occurred.

4785 **RETURN VALUE**

4786 Upon successful completion, the `atan2()`, `atan2f()`, and `atan2l()` functions shall return the arc
 4787 tangent of y/x in the range $[-\pi, \pi]$ radians. If both arguments are 0.0, an implementation-defined
 4788 value is returned and `errno` may be set to [EDOM].

4789 **XSI** If x or y is NaN, NaN shall be returned and `errno` may be set to [EDOM].

4790 If the result underflows, 0.0 shall be returned and `errno` may be set to [ERANGE].

4791 **ERRORS**

4792 The `atan2()`, `atan2f()`, and `atan2l()` functions may fail if:

4793 **XSI** [EDOM] Both arguments are 0.0 or one or more of the arguments is NaN.

4794 [ERANGE] The result underflows

4795 **XSI** No other errors shall occur.

4796 **EXAMPLES**

4797 None.

4798 **APPLICATION USAGE**

4799 None.

4800 **RATIONALE**

4801 None.

4802 **FUTURE DIRECTIONS**

4803 None.

4804 **SEE ALSO**4805 `atan()`, `isnan()`, `tan()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<math.h>`4806 **CHANGE HISTORY**

4807 First released in Issue 1. Derived from Issue 1 of the SVID.

4808 **Issue 4**4809 References to `matherr()` are removed.

4810 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the
 4811 ISO C standard and to rationalize error handling in the mathematics functions.

4812 The return value specified for [EDOM] is marked as an extension.

4813 **Issue 5**

4814 The DESCRIPTION is updated to indicate how an application should check for an error. This
4815 text was previously published in the APPLICATION USAGE section.

4816 **NAME**

4817 atanh — hyperbolic arc tangent

4818 **SYNOPSIS**

4819 xSI #include <math.h>

4820 double atanh(double x);

4821

4822 **DESCRIPTION**

4823 Refer to *acosh()*.

4824 **NAME**

4825 atexit — register a function to run at process termination

4826 **SYNOPSIS**

4827 #include <stdlib.h>

4828 int atexit(void (**func*)(void));4829 **DESCRIPTION**4830 cx The functionality described on this reference page is aligned with the ISO C standard. Any
4831 conflict between the requirements described here and the ISO C standard is unintentional. This
4832 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.4833 The *atexit()* function registers the function pointed to by *func*, to be called without arguments at
4834 normal program termination. At normal program termination, all functions registered by the
4835 cx *atexit()* function shall be called, in the reverse order of their registration. Normal termination
4836 occurs either by a call to *exit()* or a return from *main()*.4837 At least 32 functions can be registered with *atexit()*.4838 cx After a successful call to any of the *exec* functions, any functions previously registered by *atexit()*
4839 shall no longer be registered.4840 **RETURN VALUE**4841 Upon successful completion, *atexit()* shall return 0; otherwise, it shall return a non-zero value.4842 **ERRORS**

4843 No errors are defined.

4844 **EXAMPLES**

4845 None.

4846 **APPLICATION USAGE**4847 The functions registered by a call to *atexit()* must return to ensure that all registered functions
4848 are called.4849 The application should call *sysconf()* to obtain the value of {ATEXIT_MAX}, the number of
4850 functions that can be registered. There is no way for an application to tell how many functions
4851 have already been registered with *atexit()*.4852 **RATIONALE**

4853 None.

4854 **FUTURE DIRECTIONS**

4855 None.

4856 **SEE ALSO**4857 *exit()*, *sysconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>4858 **CHANGE HISTORY**

4859 First released in Issue 4. Derived from the ANSI C standard.

4860 **Issue 4, Version 2**4861 The APPLICATION USAGE section is updated to indicate how an application can determine the
4862 setting of {ATEXIT_MAX}, which is a constant added for X/OPEN UNIX conformance.4863 **Issue 6**

4864 Extensions beyond the ISO C standard are now marked.

4865 **NAME**

4866 atof — convert a string to double-precision number

4867 **SYNOPSIS**

4868 #include <stdlib.h>

4869 double atof(const char **str*);4870 **DESCRIPTION**

4871 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
4872 conflict between the requirements described here and the ISO C standard is unintentional. This
4873 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

4874 The call *atof(str)* shall be equivalent to:4875 *strtod(str, (char **)NULL)*,

4876 except that the handling of errors may differ. If the value cannot be represented, the behavior is
4877 undefined.

4878 **RETURN VALUE**4879 The *atof()* function shall return the converted value if the value can be represented.4880 **ERRORS**

4881 No errors are defined.

4882 **EXAMPLES**

4883 None.

4884 **APPLICATION USAGE**

4885 The *atof()* function is subsumed by *strtod()* but is retained because it is used extensively in
4886 existing code. If the number is not known to be in range, *strtod()* should be used because *atof()* is
4887 not required to perform any error checking.

4888 **RATIONALE**

4889 None.

4890 **FUTURE DIRECTIONS**

4891 None.

4892 **SEE ALSO**4893 *strtod()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>4894 **CHANGE HISTORY**

4895 First released in Issue 1. Derived from Issue 1 of the SVID.

4896 **Issue 4**4897 Reference to how *str* is converted is removed from the DESCRIPTION.

4898 The APPLICATION USAGE section is added.

4899 The following change is incorporated for alignment with the ISO C standard:

- 4900
- The type of argument *str* is changed from **char*** to **const char***.

4901 **NAME**

4902 atoi — convert a string to an integer

4903 **SYNOPSIS**

4904 #include <stdlib.h>

4905 int atoi(const char *str);

4906 **DESCRIPTION**

4907 cx The functionality described on this reference page is aligned with the ISO C standard. Any
4908 conflict between the requirements described here and the ISO C standard is unintentional. This
4909 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

4910 The call *atoi(str)* shall be equivalent to:

4911 (int) strtol(str, (char **)NULL, 10)

4912 except that the handling of errors may differ. If the value cannot be represented, the behavior is
4913 undefined.4914 **RETURN VALUE**4915 The *atoi()* function shall return the converted value if the value can be represented.4916 **ERRORS**

4917 No errors are defined.

4918 **EXAMPLES**4919 **Converting an Argument**

4920 The following example checks for proper usage of the program. If there is an argument and the
4921 decimal conversion of this argument (obtained using *atoi()*) is greater than 0, then the program
4922 has a valid number of minutes to wait for an event.

```
4923         #include <stdlib.h>
4924         #include <stdio.h>
4925         ...
4926         int minutes_to_event;
4927         ...
4928         if (argc < 2 || ((minutes_to_event = atoi (argv[1]))) <= 0) {
4929             fprintf(stderr, "Usage: %s minutes\n", argv[0]); exit(1);
4930         }
4931         ...
```

4932 **APPLICATION USAGE**

4933 The *atoi()* function is subsumed by *strtol()* but is retained because it is used extensively in
4934 existing code. If the number is not known to be in range, *strtol()* should be used because *atoi()* is
4935 not required to perform any error checking.

4936 **RATIONALE**

4937 None.

4938 **FUTURE DIRECTIONS**

4939 None.

4940 **SEE ALSO**4941 *strtol()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

4942 **CHANGE HISTORY**

4943 First released in Issue 1. Derived from Issue 1 of the SVID.

4944 **Issue 4**

4945 Reference to how *str* is converted is removed from the DESCRIPTION.

4946 The APPLICATION USAGE section is added.

4947 The following change is incorporated for alignment with the ISO C standard:

- 4948 • The type of argument *str* is changed from **char*** to **const char***.

4949 **NAME**4950 `atol, atoll` — convert a string to a long integer4951 **SYNOPSIS**4952 `#include <stdlib.h>`4953 `long atol(const char *str);`4954 `long long atoll(const char *nptr);`4955 **DESCRIPTION**

4956 **cx** The functionality described on this reference page is aligned with the ISO C standard. Any
4957 conflict between the requirements described here and the ISO C standard is unintentional. This
4958 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

4959 The call `atol(str)` shall be equivalent to:4960 `strtoul(str, (char **)NULL, 10)`4961 The call `atoll(str)` shall be equivalent to:4962 `strtoll(nptr, (char **)NULL, 10)`

4963 except that the handling of errors may differ. If the value cannot be represented, the behavior is
4964 undefined.

4965 **RETURN VALUE**

4966 These functions shall return the converted value if the value can be represented.

4967 **ERRORS**

4968 No errors are defined.

4969 **EXAMPLES**

4970 None.

4971 **APPLICATION USAGE**

4972 The `atol()` function is subsumed by `strtoul()` but is retained because it is used extensively in
4973 existing code. If the number is not known to be in range, `strtoul()` should be used because `atol()` is
4974 not required to perform any error checking.

4975 **RATIONALE**

4976 None.

4977 **FUTURE DIRECTIONS**

4978 None.

4979 **SEE ALSO**4980 `strtoul()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>`4981 **CHANGE HISTORY**

4982 First released in Issue 1. Derived from Issue 1 of the SVID.

4983 **Issue 4**4984 Reference to how `str` is converted is removed from the DESCRIPTION.

4985 The APPLICATION USAGE section is added.

4986 The following changes are incorporated for alignment with the ISO C standard:

- 4987 • The type of argument `str` is changed from `char*` to `const char*`.
- 4988 • The return type of the function is expanded to `long`.

4989 **Issue 6**
4990

The *atoll()* function is added for alignment with the ISO/IEC 9899:1999 standard.

4991 **NAME**4992 **basename** — return the last component of a path name4993 **SYNOPSIS**4994 **XSI** #include <libgen.h>

4995 char *basename(char *path);

4996

4997 **DESCRIPTION**4998 The *basename()* function shall take the path name pointed to by *path* and return a pointer to the
4999 final component of the path name, deleting any trailing '/' characters.5000 If the string consists entirely of the '/' character, *basename()* shall return a pointer to the string
5001 "/". If the string is exactly "///", it is implementation-defined whether '/' or "///" is
5002 returned.5003 If *path* is a null pointer or points to an empty string, *basename()* shall return a pointer to the
5004 string ".".5005 The *basename()* function may modify the string pointed to by *path*, and may return a pointer to
5006 static storage that may then be overwritten by a subsequent call to *basename()*.5007 The *basename()* function need not be reentrant. A function that is not required to be reentrant is
5008 not required to be thread-safe.5009 **RETURN VALUE**5010 The *basename()* function shall return a pointer to the final component of *path*.5011 **ERRORS**

5012 No errors are defined.

5013 **EXAMPLES**5014 **Using basename()**5015 The following program fragment returns a pointer to the value *lib*, which is the base name of
5016 */usr/lib*.5017 #include <libgen.h>
5018 ...
5019 char *name = "/usr/lib";
5020 char *base;

5021 base = basename(name);
5022 ...5023 **Sample Input and Output Strings for basename()**5024 In the following table, the input string is the value pointed to by *path*, and the output string is
5025 the return value of the *basename()* function.

5026

5027

5028

5029

Input String	Output String
"/usr/lib"	"lib"
"/usr/"	"usr"
"/"	"/"

5030 **APPLICATION USAGE**

5031 None.

5032 **RATIONALE**

5033 None.

5034 **FUTURE DIRECTIONS**

5035 None.

5036 **SEE ALSO**5037 *dirname()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<libgen.h>**, the Shell and
5038 Utilities volume of IEEE Std. 1003.1-200x, *basename*5039 **CHANGE HISTORY**

5040 First released in Issue 4, Version 2.

5041 **Issue 5**

5042 Moved from X/OPEN UNIX extension to BASE.

5043 Normative text previously in the APPLICATION USAGE section is moved to the
5044 DESCRIPTION.

5045 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

5046 **Issue 6**

5047 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

5048 **NAME**5049 bcmp — memory operations (**LEGACY**)5050 **SYNOPSIS**

5051 xSI #include <strings.h>

5052 int bcmp(const void *s1, const void *s2, size_t n);

5053

5054 **DESCRIPTION**5055 The *bcmp()* function shall compare the first *n* bytes of the area pointed to by *s1* with the area pointed to by *s2*.5057 **RETURN VALUE**5058 The *bcmp()* function shall return 0 if *s1* and *s2* are identical; otherwise, it shall return non-zero.
5059 Both areas are assumed to be *n* bytes long. If the value of *n* is 0, *bcmp()* shall return 0.5060 **ERRORS**

5061 No errors are defined.

5062 **EXAMPLES**

5063 None.

5064 **APPLICATION USAGE**5065 *memcmp()* is preferred over this function.5066 For maximum portability, it is recommended to replace the function call to *bcmp()* as follows:

5067 #define bcmp(b1,b2,len) memcmp((b1), (b2), (size_t)(len))

5068 **RATIONALE**

5069 None.

5070 **FUTURE DIRECTIONS**

5071 This function may be withdrawn in a future version.

5072 **SEE ALSO**5073 *memcmp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <strings.h>5074 **CHANGE HISTORY**

5075 First released in Issue 4, Version 2.

5076 **Issue 5**

5077 Moved from X/OPEN UNIX extension to BASE.

5078 **Issue 6**

5079 This function is marked LEGACY.

5080 **NAME**5081 `bcopy` — memory operations (**LEGACY**)5082 **SYNOPSIS**5083 XSI

```
#include <strings.h>
```

5084

```
void bcopy(const void *s1, void *s2, size_t n);
```

5085

5086 **DESCRIPTION**5087 The `bcopy()` function shall copy `n` bytes from the area pointed to by `s1` to the area pointed to by
5088 `s2`.5089 The bytes are copied correctly even if the area pointed to by `s1` overlaps the area pointed to by
5090 `s2`.5091 **RETURN VALUE**5092 The `bcopy()` function shall return no value.5093 **ERRORS**

5094 No errors are defined.

5095 **EXAMPLES**

5096 None.

5097 **APPLICATION USAGE**5098 `memmove()` is preferred over this function.

5099 The following are approximately equivalent (note the order of the arguments):

5100 `bcopy(s1, s2, n) ~ memmove(s2, s1, n)`5101 For maximum portability, it is recommended to replace the function call to `bcopy()` as follows:5102

```
#define bcopy(b1,b2,len) (memmove((b2), (b1), (len)), (void) 0)
```

5103 **RATIONALE**

5104 None.

5105 **FUTURE DIRECTIONS**

5106 This function may be withdrawn in a future version.

5107 **SEE ALSO**5108 `memmove()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<strings.h>`5109 **CHANGE HISTORY**

5110 First released in Issue 4, Version 2.

5111 **Issue 5**

5112 Moved from X/OPEN UNIX extension to BASE.

5113 **Issue 6**

5114 This function is marked LEGACY.

5115 **NAME**

5116 bind — bind a name to a socket

5117 **SYNOPSIS**

5118 #include <sys/socket.h>

5119 int bind(int *socket*, const struct sockaddr **address*,
5120 socklen_t *address_len*);5121 **DESCRIPTION**5122 The *bind()* function shall assign a local socket address *address* to a socket identified by descriptor
5123 *socket* that has no local socket address assigned. Sockets created with the *socket()* function are
5124 initially unnamed; they are identified only by their address family.5125 The *bind()* function takes the following arguments:5126 *socket* Specifies the file descriptor of the socket to be bound.5127 *address* Points to a **sockaddr** structure containing the address to be bound to the
5128 socket. The length and format of the address depend on the address family of
5129 the socket.5130 *address_len* Specifies the length of the **sockaddr** structure pointed to by the *address*
5131 argument.5132 The socket specified by *socket* may require the process to have appropriate privileges to use the
5133 *bind()* function.5134 **RETURN VALUE**5135 Upon successful completion, *bind()* shall return 0; otherwise, -1 shall be returned and *errno* set
5136 to indicate the error.5137 **ERRORS**5138 The *bind()* function shall fail if:

5139 [EADDRINUSE]

5140 The specified address is already in use.

5141 [EADDRNOTAVAIL]

5142 The specified address is not available from the local machine.

5143 [EAFNOSUPPORT]

5144 The specified address is not a valid address for the address family of the
5145 specified socket.5146 [EBADF] The *socket* argument is not a valid file descriptor.5147 [EINVAL] The socket is already bound to an address, and the protocol does not support
5148 binding to a new address; or the socket has been shut down.5149 [ENOTSOCK] The *socket* argument does not refer to a socket.5150 [EOPNOTSUPP] The socket type of the specified socket does not support binding to an
5151 address.5152 If the address family of the socket is AF_UNIX, then *bind()* shall fail if:5153 [EACCES] A component of the path prefix denies search permission, or the requested
5154 name requires writing in a directory with a mode that denies write
5155 permission.

5156	[EDESTADDRREQ] or [EISDIR]	
5157		The <i>address</i> argument is a null pointer.
5158	[EIO]	An I/O error occurred.
5159	[ELOOP]	A loop exists in symbolic links encountered during resolution of the path name in <i>address</i> .
5160		
5161	[ENAMETOOLONG]	
5162		A component of a path name exceeded {NAME_MAX} characters, or an entire path name exceeded {PATH_MAX} characters.
5163		
5164	[ENOENT]	A component of the path name does not name an existing file or the path name is an empty string.
5165		
5166	[ENOTDIR]	A component of the path prefix of the path name in <i>address</i> is not a directory.
5167	[EROFS]	The name would reside on a read-only file system.
5168		The <i>bind()</i> function may fail if:
5169	[EACCES]	The specified address is protected and the current user does not have permission to bind to it.
5170		
5171	[EINVAL]	The <i>address_len</i> argument is not a valid length for the address family.
5172	[EISCONN]	The socket is already connected.
5173	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the path name in <i>address</i> .
5174		
5175	[ENAMETOOLONG]	
5176		Path name resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.
5177		
5178	[ENOBUFS]	Insufficient resources were available to complete the call.
5179	EXAMPLES	
5180		None.
5181	APPLICATION USAGE	
5182		An application program can retrieve the assigned socket name with the <i>getsockname()</i> function.
5183	RATIONALE	
5184		None.
5185	FUTURE DIRECTIONS	
5186		None.
5187	SEE ALSO	
5188		<i>connect()</i> , <i>getsockname()</i> , <i>listen()</i> , <i>socket()</i> , the Base Definitions volume of IEEE Std. 1003.1-200x,
5189		< sys/socket.h >
5190	CHANGE HISTORY	
5191		First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

5192 **NAME**5193 **bsd_signal** — simplified signal facilities5194 **SYNOPSIS**5195 **OB** #include <signal.h>5196 void (*bsd_signal(int *sig*, void (**func*)(int)))(int);

5197

5198 **DESCRIPTION**5199 The *bsd_signal()* function provides a partially compatible interface for programs written to
5200 historical system interfaces (see APPLICATION USAGE).5201 The function call *bsd_signal(sig, func)* has an effect as if implemented as:5202 void (*bsd_signal(int *sig*, void (**func*)(int)))(int)

5203 {

5204 struct sigaction *act*, *oact*;5205 *act*.sa_handler = *func*;5206 *act*.sa_flags = SA_RESTART;5207 sigemptyset(&*act*.sa_mask);5208 sigaddset(&*act*.sa_mask, *sig*);5209 if (sigaction(*sig*, &*act*, &*oact*) == -1)

5210 return(SIG_ERR);

5211 return(*oact*.sa_handler);

5212 }

5213 The handler function should be declared:

5214 void handler(int *sig*);5215 where *sig* is the signal number. The behavior is undefined if *func* is a function that takes more
5216 than one argument, or an argument of a different type.5217 **RETURN VALUE**5218 Upon successful completion, *bsd_signal()* shall return the previous action for *sig*. Otherwise,
5219 SIG_ERR shall be returned and *errno* shall be set to indicate the error.5220 **ERRORS**5221 Refer to *sigaction()*.5222 **EXAMPLES**

5223 None.

5224 **APPLICATION USAGE**5225 This function is a direct replacement for the BSD *signal()* function for simple applications that
5226 are installing a single-argument signal handler function. If a BSD signal handler function is being
5227 installed that expects more than one argument, the application has to be modified to use
5228 *sigaction()*. The *bsd_signal()* function differs from *signal()* in that the SA_RESTART flag is set
5229 and the SA_RESETHAND is clear when *bsd_signal()* is used. The state of these flags is not
5230 specified for *signal()*.5231 It is recommended that new applications use the *sigaction()* function.5232 **RATIONALE**

5233 None.

5234 **FUTURE DIRECTIONS**

5235 None.

5236 **SEE ALSO**5237 *sigaction()*, *sigaddset()*, *sigemptyset()*, *signal()*, the Base Definitions volume of
5238 IEEE Std. 1003.1-200x, <**signal.h**>5239 **CHANGE HISTORY**

5240 First released in Issue 4, Version 2.

5241 **Issue 5**

5242 Moved from X/OPEN UNIX extension to BASE.

5243 **Issue 6**

5244 This function is marked obsolescent.

5245 **NAME**5246 **bsearch** — binary search a sorted table5247 **SYNOPSIS**

5248 #include <stdlib.h>

5249 void *bsearch(const void *key, const void *base, size_t nel,
5250 size_t width, int (*compar)(const void *, const void *));5251 **DESCRIPTION**5252 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5253 conflict between the requirements described here and the ISO C standard is unintentional. This
5254 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.5255 The *bsearch()* function searches an array of *nel* objects, the initial element of which is pointed to
5256 by *base*, for an element that matches the object pointed to by *key*. The size of each element in the
5257 array is specified by *width*.5258 The comparison function pointed to by *compar* is called with two arguments that point to the *key*
5259 object and to an array element, in that order.5260 The application shall ensure that the function returns an integer less than, equal to, or greater
5261 than 0 if the *key* object is considered, respectively, to be less than, to match, or to be greater than
5262 the array element. The application shall ensure that the array consists of all the elements that
5263 compare less than, all the elements that compare equal to, and all the elements that compare
5264 greater than the *key* object, in that order.5265 **RETURN VALUE**5266 The *bsearch()* function shall return a pointer to a matching member of the array, or a null pointer
5267 if no match is found. If two or more members compare equal, which member is returned is
5268 unspecified.5269 **ERRORS**

5270 No errors are defined.

5271 **EXAMPLES**5272 The example below searches a table containing pointers to nodes consisting of a string and its
5273 length. The table is ordered alphabetically on the string in the node pointed to by each entry.5274 The code fragment below reads in strings and either finds the corresponding node and prints out
5275 the string and its length, or prints an error message.5276 #include <stdio.h>
5277 #include <stdlib.h>
5278 #include <string.h>

5279 #define TABSIZE 1000

5280 struct node { /* These are stored in the table. */
5281 char *string;
5282 int length;
5283 };
5284 struct node table[TABSIZE]; /* Table to be searched. */
5285 .
5286 .
5287 .
5288 {
5289 struct node *node_ptr, node;
5290 /* routine to compare 2 nodes */

```

5291     int node_compare(const void *, const void *);
5292     char str_space[20]; /* Space to read string into. */
5293     .
5294     .
5295     .
5296     node.string = str_space;
5297     while (scanf("%s", node.string) != EOF) {
5298         node_ptr = (struct node *)bsearch((void *)&node,
5299             (void *)table, TABSIZE,
5300             sizeof(struct node), node_compare);
5301         if (node_ptr != NULL) {
5302             (void)printf("string = %20s, length = %d\n",
5303                 node_ptr->string, node_ptr->length);
5304         } else {
5305             (void)printf("not found: %s\n", node.string);
5306         }
5307     }
5308 }
5309 /*
5310     This routine compares two nodes based on an
5311     alphabetical ordering of the string field.
5312 */
5313 int
5314 node_compare(const void *node1, const void *node2)
5315 {
5316     return strcoll(((const struct node *)node1)->string,
5317         ((const struct node *)node2)->string);
5318 }

```

5319 APPLICATION USAGE

5320 The pointers to the key and the element at the base of the table should be of type pointer-to-
5321 element.

5322 The comparison function need not compare every byte, so arbitrary data may be contained in
5323 the elements in addition to the values being compared.

5324 In practice, the array is usually sorted according to the comparison function.

5325 RATIONALE

5326 None.

5327 FUTURE DIRECTIONS

5328 None.

5329 SEE ALSO

5330 *hcreate()*, *lsearch()*, *qsort()*, *tsearch()*, the Base Definitions volume of IEEE Std. 1003.1-200x,
5331 `<stdlib.h>`

5332 CHANGE HISTORY

5333 First released in Issue 1. Derived from Issue 1 of the SVID.

5334 Issue 4

5335 Text indicating the need for various casts is removed from the APPLICATION USAGE section.

5336 The code in the EXAMPLES section is changed to use *strcoll()* instead of *strcmp()* in
5337 *node_compare()*.

- 5338 The return value and the contents of the array are now requirements on the application.
- 5339 The DESCRIPTION is changed to specify the order of arguments.
- 5340 The following changes are incorporated for alignment with the ISO C standard:
- 5341 • The type of arguments *key* and *base*, and the type of arguments to *compar*, are changed from
 - 5342 **void*** to **const void***.
 - 5343 • The requirement that the table be sorted according to *compar* is removed from the
 - 5344 DESCRIPTION.
- 5345 **Issue 6**
- 5346 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

5347 **NAME**

5348 btowc — single-byte to wide-character conversion

5349 **SYNOPSIS**

5350 #include <stdio.h>

5351 #include <wchar.h>

5352 wint_t btowc(int c);

5353 **DESCRIPTION**

5354 cx The functionality described on this reference page is aligned with the ISO C standard. Any
5355 conflict between the requirements described here and the ISO C standard is unintentional. This
5356 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

5357 The *btowc()* function shall determine whether *c* constitutes a valid (one-byte) character in the
5358 initial shift state.

5359 The behavior of this function shall be affected by the *LC_CTYPE* category of the current locale.

5360 **RETURN VALUE**

5361 The *btowc()* function shall return WEOF if *c* has the value EOF or if (**unsigned char**) *c* does not
5362 constitute a valid (one-byte) character in the initial shift state. Otherwise, it shall return the
5363 wide-character representation of that character.

5364 **ERRORS**

5365 No errors are defined.

5366 **EXAMPLES**

5367 None.

5368 **APPLICATION USAGE**

5369 None.

5370 **RATIONALE**

5371 None.

5372 **FUTURE DIRECTIONS**

5373 None.

5374 **SEE ALSO**5375 *wctob()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <wchar.h>5376 **CHANGE HISTORY**

5377 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995
5378 (E).

5379 **NAME**5380 bzero — memory operations (**LEGACY**)5381 **SYNOPSIS**

5382 XSI #include <strings.h>

5383 void bzero(void *s, size_t n);

5384

5385 **DESCRIPTION**5386 The *bzero()* function shall place *n* zero-valued bytes in the area pointed to by *s*.5387 **RETURN VALUE**5388 The *bzero()* function shall return no value.5389 **ERRORS**

5390 No errors are defined.

5391 **EXAMPLES**

5392 None.

5393 **APPLICATION USAGE**5394 *memset()* is preferred over this function.5395 For maximum portability, it is recommended to replace the function call to *bzero()* as follows:

5396 #define bzero(b,len) (memset((b), '\0', (len)), (void) 0)

5397 **RATIONALE**

5398 None.

5399 **FUTURE DIRECTIONS**

5400 This function may be withdrawn in a future version.

5401 **SEE ALSO**5402 *memset()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <strings.h>5403 **CHANGE HISTORY**

5404 First released in Issue 4, Version 2.

5405 **Issue 5**

5406 Moved from X/OPEN UNIX extension to BASE.

5407 **Issue 6**

5408 This function is marked LEGACY.

5409 **NAME**

5410 cabs, cabsf, cabsl — return a complex absolute value

5411 **SYNOPSIS**

5412 #include <complex.h>

5413 double cabs(double complex *z*);5414 float cabsf(float complex *z*);5415 long double cabsl(long double complex *z*);5416 **DESCRIPTION**

5417 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5418 conflict between the requirements described here and the ISO C standard is unintentional. This
5419 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

5420 These functions shall compute the complex absolute value (also called norm, modulus, or
5421 magnitude) of *z*.

5422 **RETURN VALUE**

5423 These functions shall return the complex absolute value.

5424 **ERRORS**

5425 No errors are defined.

5426 **EXAMPLES**

5427 None.

5428 **APPLICATION USAGE**

5429 None.

5430 **RATIONALE**

5431 None.

5432 **FUTURE DIRECTIONS**

5433 None.

5434 **SEE ALSO**

5435 The Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>

5436 **CHANGE HISTORY**

5437 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5438 **NAME**5439 `acos`, `acosf`, `acosl` — complex arc cosine functions5440 **SYNOPSIS**5441 `#include <complex.h>`5442 `double complex cacos(double complex z);`5443 `float complex cacosf(float complex z);`5444 `long double complex cacosl(long double complex z);`5445 **DESCRIPTION**5446 `CX` The functionality described on this reference page is aligned with the ISO C standard. Any
5447 conflict between the requirements described here and the ISO C standard is unintentional. This
5448 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.5449 These functions shall compute the complex arc cosine of z , with branch cuts outside the interval
5450 $[-1, +1]$ along the real axis.5451 **RETURN VALUE**5452 These functions shall return the complex arc cosine value, in the range of a strip mathematically
5453 unbounded along the imaginary axis and in the interval $[0, \pi]$ along the real axis.5454 **ERRORS**

5455 No errors are defined.

5456 **EXAMPLES**

5457 None.

5458 **APPLICATION USAGE**

5459 None.

5460 **RATIONALE**

5461 None.

5462 **FUTURE DIRECTIONS**

5463 None.

5464 **SEE ALSO**5465 `ccos()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<complex.h>`5466 **CHANGE HISTORY**

5467 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5468 **NAME**
5469 cacosh, cacoshf, cacoshl — complex arc hyperbolic cosine functions

5470 **SYNOPSIS**
5471 #include <complex.h>
5472 double complex cacosh(double complex z);
5473 float complex cacoshf(float complex z);
5474 long double complex cacoshl(long double complex z);

5475 **DESCRIPTION**
5476 cx The functionality described on this reference page is aligned with the ISO C standard. Any
5477 conflict between the requirements described here and the ISO C standard is unintentional. This
5478 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.
5479 These functions shall compute the complex arc hyperbolic cosine of z , with a branch cut at
5480 values less than 1 along the real axis.

5481 **RETURN VALUE**
5482 These functions shall return the complex arc hyperbolic cosine value, in the range of a half-strip
5483 of non-negative values along the real axis and in the interval $[-i\pi, +i\pi]$ along the imaginary axis.

5484 **ERRORS**
5485 No errors are defined.

5486 **EXAMPLES**
5487 None.

5488 **APPLICATION USAGE**
5489 None.

5490 **RATIONALE**
5491 None.

5492 **FUTURE DIRECTIONS**
5493 None.

5494 **SEE ALSO**
5495 *ccosh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>

5496 **CHANGE HISTORY**
5497 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5498 **NAME**

5499 calloc — a memory allocator

5500 **SYNOPSIS**

5501 #include <stdlib.h>

5502 void *calloc(size_t *nelem*, size_t *elsize*);5503 **DESCRIPTION**5504 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5505 conflict between the requirements described here and the ISO C standard is unintentional. This
5506 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.5507 The *calloc()* function allocates unused space for an array of *nelem* elements each of whose size in
5508 bytes is *elsize*. The space is initialized to all bits 0.5509 The order and contiguity of storage allocated by successive calls to *calloc()* is unspecified. The
5510 pointer returned if the allocation succeeds is suitably aligned so that it may be assigned to a
5511 pointer to any type of object and then used to access such an object or an array of such objects in
5512 the space allocated (until the space is explicitly freed or reallocated). Each such allocation shall
5513 yield a pointer to an object disjoint from any other object. The pointer returned points to the start
5514 (lowest byte address) of the allocated space. If the space cannot be allocated, a null pointer is
5515 returned. If the size of the space requested is 0, the behavior is implementation-defined; the
5516 value returned shall be either a null pointer or a unique pointer.5517 **RETURN VALUE**5518 Upon successful completion with both *nelem* and *elsize* non-zero, *calloc()* shall return a pointer to
5519 the allocated space. If either *nelem* or *elsize* is 0, then either a null pointer or a unique pointer
5520 value that can be successfully passed to *free()* shall be returned. Otherwise, it shall return a null
5521 **CX** pointer and set *errno* to indicate the error.5522 **ERRORS**5523 The *calloc()* function shall fail if:5524 **CX** [ENOMEM] Insufficient memory is available.5525 **EXAMPLES**

5526 None.

5527 **APPLICATION USAGE**

5528 There is now no requirement for the implementation to support the inclusion of <malloc.h>.

5529 **RATIONALE**

5530 None.

5531 **FUTURE DIRECTIONS**

5532 None.

5533 **SEE ALSO**5534 *free()*, *malloc()*, *realloc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>5535 **CHANGE HISTORY**

5536 First released in Issue 1. Derived from Issue 1 of the SVID.

5537 **Issue 4**5538 The setting of *errno* and the [ENOMEM] error are marked as extensions.5539 The APPLICATION USAGE section is changed to record that <malloc.h> need no longer be
5540 supported on XSI-conformant systems.

- 5541 The following changes are incorporated in this issue for alignment with the ISO C standard:
- 5542 • The DESCRIPTION is updated to indicate:
- 5543 — The order and contiguity of storage allocated by successive calls to this function is
5544 unspecified.
- 5545 — Each allocation yields a pointer to an object disjoint from any other object.
- 5546 — The returned pointer points to the lowest byte address of the allocation.
- 5547 — The behavior if space is requested of zero size.
- 5548 • The RETURN VALUE section is updated to indicate what is returned if either *nelem* or *elsize*
5549 is 0.

5550 Issue 6

5551 Extensions beyond the ISO C standard are now marked.

5552 The following new requirements on POSIX implementations derive from alignment with the
5553 Single UNIX Specification:

- 5554 • The setting of *errno* and the [ENOMEM] error condition are mandatory if an insufficient
5555 memory condition occurs.

5556 **NAME**

5557 carg, cargf, cargl — complex argument functions

5558 **SYNOPSIS**

5559 #include <complex.h>

5560 double carg(double complex *z*);5561 float cargf(float complex *z*);5562 long double cargl(long double complex *z*);5563 **DESCRIPTION**

5564 cx The functionality described on this reference page is aligned with the ISO C standard. Any
5565 conflict between the requirements described here and the ISO C standard is unintentional. This
5566 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

5567 These functions shall compute the argument (also called phase angle) of *z*, with a branch cut
5568 along the negative real axis.

5569 **RETURN VALUE**5570 These functions shall return the value of the argument in the interval $[-\pi, +\pi]$.5571 **ERRORS**

5572 No errors are defined.

5573 **EXAMPLES**

5574 None.

5575 **APPLICATION USAGE**

5576 None.

5577 **RATIONALE**

5578 None.

5579 **FUTURE DIRECTIONS**

5580 None.

5581 **SEE ALSO**5582 *cimag()*, *conj()*, *cproj()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>5583 **CHANGE HISTORY**

5584 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5585 **NAME**

5586 casin, casinf, casinl — complex arc sine functions

5587 **SYNOPSIS**

5588 #include <complex.h>

5589 double complex casin(double complex *z*);5590 float complex casinf(float complex *z*);5591 long double complex casinl(long double complex *z*);5592 **DESCRIPTION**5593 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5594 conflict between the requirements described here and the ISO C standard is unintentional. This
5595 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.5596 These functions shall compute the complex arc sine of *z*, with branch cuts outside the interval
5597 $[-1, +1]$ along the real axis.5598 **RETURN VALUE**5599 These functions shall return the complex arc sine value, in the range of a strip mathematically
5600 unbounded along the imaginary axis and in the interval $[-\pi/2, +\pi/2]$ along the real axis.5601 **ERRORS**

5602 No errors are defined.

5603 **EXAMPLES**

5604 None.

5605 **APPLICATION USAGE**

5606 None.

5607 **RATIONALE**

5608 None.

5609 **FUTURE DIRECTIONS**

5610 None.

5611 **SEE ALSO**5612 `csin()`, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>5613 **CHANGE HISTORY**

5614 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5615 **NAME**

5616 casinh, casinhf, casinhl — complex arc hyperbolic sine functions

5617 **SYNOPSIS**

5618 #include <complex.h>

5619 double complex casinh(double complex *z*);5620 float complex casinhf(float complex *z*);5621 long double complex casinhl(long double complex *z*);5622 **DESCRIPTION**5623 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5624 conflict between the requirements described here and the ISO C standard is unintentional. This
5625 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.5626 These functions shall compute the complex arc hyperbolic sine of *z*, with branch cuts outside the
5627 interval $[-i, +i]$ along the imaginary axis.5628 **RETURN VALUE**5629 These functions shall return the complex arc hyperbolic sine value, in the range of a strip
5630 mathematically unbounded along the real axis and in the interval $[-i\pi/2, +i\pi/2]$ along the
5631 imaginary axis.5632 **ERRORS**

5633 No errors are defined.

5634 **EXAMPLES**

5635 None.

5636 **APPLICATION USAGE**

5637 None.

5638 **RATIONALE**

5639 None.

5640 **FUTURE DIRECTIONS**

5641 None.

5642 **SEE ALSO**5643 `csinh()`, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>5644 **CHANGE HISTORY**

5645 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5646 **NAME**

5647 catan, catanf, catanl — complex arc tangent functions

5648 **SYNOPSIS**

5649 #include <complex.h>

5650 double complex catan(double complex *z*);5651 float complex catanf(float complex *z*);5652 long double complex catanl(long double complex *z*);5653 **DESCRIPTION**5654 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5655 conflict between the requirements described here and the ISO C standard is unintentional. This
5656 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.5657 These functions shall compute the complex arc tangent of *z*, with branch cuts outside the
5658 interval $[-i, +i]$ along the imaginary axis.5659 **RETURN VALUE**5660 These functions shall return the complex arc tangent value, in the range of a strip
5661 mathematically unbounded along the imaginary axis and in the interval $[-\pi/2, +\pi/2]$ along the
5662 real axis.5663 **ERRORS**

5664 No errors are defined.

5665 **EXAMPLES**

5666 None.

5667 **APPLICATION USAGE**

5668 None.

5669 **RATIONALE**

5670 None.

5671 **FUTURE DIRECTIONS**

5672 None.

5673 **SEE ALSO**

5674 ctan(), the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>

5675 **CHANGE HISTORY**

5676 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5677 **NAME**

5678 catanh, catanhf, catanhl — complex arc hyperbolic tangent functions

5679 **SYNOPSIS**

5680 #include <complex.h>

5681 double complex catanh(double complex *z*);5682 float complex catanhf(float complex *z*);5683 long double complex catanhl(long double complex *z*);5684 **DESCRIPTION**

5685 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5686 conflict between the requirements described here and the ISO C standard is unintentional. This
5687 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

5688 These functions shall compute the complex arc hyperbolic tangent of *z*, with branch cuts outside
5689 the interval $[-1, +1]$ along the real axis.

5690 **RETURN VALUE**

5691 These functions shall return the complex arc hyperbolic tangent value, in the range of a strip
5692 mathematically unbounded along the real axis and in the interval $[-i\pi/2, +i\pi/2]$ along the
5693 imaginary axis.

5694 **ERRORS**

5695 No errors are defined.

5696 **EXAMPLES**

5697 None.

5698 **APPLICATION USAGE**

5699 None.

5700 **RATIONALE**

5701 None.

5702 **FUTURE DIRECTIONS**

5703 None.

5704 **SEE ALSO**5705 *ctanh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>5706 **CHANGE HISTORY**

5707 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5708 **NAME**

5709 catclose — close a message catalog descriptor

5710 **SYNOPSIS**

5711 xSI #include <nl_types.h>

5712 int catclose(nl_catd catd);

5713

5714 **DESCRIPTION**5715 The *catclose()* function shall close the message catalog identified by *catd*. If a file descriptor is
5716 used to implement the type **nl_catd**, that file descriptor shall be closed.5717 **RETURN VALUE**5718 Upon successful completion, *catclose()* shall return 0; otherwise, -1 shall be returned, and *errno*
5719 set to indicate the error.5720 **ERRORS**5721 The *catclose()* function may fail if:

5722 [EBADF] The catalog descriptor is not valid. |

5723 [EINTR] The *catclose()* function was interrupted by a signal. |5724 **EXAMPLES**

5725 None.

5726 **APPLICATION USAGE**

5727 None.

5728 **RATIONALE**

5729 None.

5730 **FUTURE DIRECTIONS**

5731 None.

5732 **SEE ALSO**5733 *catgets()*, *catopen()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**nl_types.h**> |5734 **CHANGE HISTORY**

5735 First released in Issue 2.

5736 **Issue 4**

5737 The [EBADF] and [EINTR] errors are added to the ERRORS section.

5738 **NAME**

5739 catgets — read a program message

5740 **SYNOPSIS**

5741 XSI #include <nl_types.h>

5742 char *catgets(nl_catd *catd*, int *set_id*, int *msg_id*, const char **s*);

5743

5744 **DESCRIPTION**

5745 The *catgets()* function attempts to read message *msg_id*, in set *set_id*, from the message catalog identified by *catd*. The *catd* argument is a message catalog descriptor returned from an earlier call to *catopen()*. The *s* argument points to a default message string which shall be returned by *catgets()* if it cannot retrieve the identified message.

5749 The *catgets()* function need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

5751 **RETURN VALUE**

5752 If the identified message is retrieved successfully, *catgets()* shall return a pointer to an internal buffer area containing the null-terminated message string. If the call is unsuccessful for any reason, *s* shall be returned and *errno* may be set to indicate the error.

5755 **ERRORS**5756 The *catgets()* function may fail if:5757 [EBADF] The *catd* argument is not a valid message catalog descriptor open for reading. |5758 [EINTR] The read operation was terminated due to the receipt of a signal, and no data |
5759 was transferred.5760 [EINVAL] The message catalog identified by *catd* is corrupted. |5761 [ENOMSG] The message identified by *set_id* and *msg_id* is not in the message catalog. |5762 **EXAMPLES**

5763 None.

5764 **APPLICATION USAGE**

5765 None.

5766 **RATIONALE**

5767 None.

5768 **FUTURE DIRECTIONS**

5769 None.

5770 **SEE ALSO**5771 *catclose()*, *catopen()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <nl_types.h> |5772 **CHANGE HISTORY**

5773 First released in Issue 2.

5774 **Issue 4**5775 The type of argument *s* is changed from **char*** to **const char***.

5776 The [EBADF] and [EINTR] errors are added to the ERRORS section.

5777 **Issue 4, Version 2**

5778 The following changes are incorporated for X/OPEN UNIX conformance:

- 5779 • The RETURN VALUE section notes that *errno* may be set to indicate an error.
- 5780 • In the ERRORS section, [EINVAL] and [ENOMSG] are added as optional errors.

5781 **Issue 5**

5782 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

5783 **Issue 6**

5784 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

5785 **NAME**5786 `catopen` — open a message catalog5787 **SYNOPSIS**5788 xSI `#include <nl_types.h>`5789 `nl_catd catopen(const char *name, int oflag);`

5790

5791 **DESCRIPTION**

5792 The `catopen()` function shall open a message catalog and return a message catalog descriptor.
 5793 The `name` argument specifies the name of the message catalog to be opened. If `name` contains a
 5794 `'/'`, then `name` specifies a complete name for the message catalog. Otherwise, the environment
 5795 variable `NLSPATH` is used with `name` substituted for `%N` (see the Base Definitions volume of
 5796 IEEE Std. 1003.1-200x, Chapter 8, Environment Variables). If `NLSPATH` does not exist in the
 5797 environment, or if a message catalog cannot be found in any of the components specified by
 5798 `NLSPATH`, then an implementation-defined default path is used. This default may be affected by
 5799 the setting of `LC_MESSAGES` if the value of `oflag` is `NL_CAT_LOCALE`, or the `LANG`
 5800 environment variable if `oflag` is 0.

5801 A message catalog descriptor remains valid in a process until that process closes it, or a
 5802 successful call to one of the `exec` functions. A change in the setting of the `LC_MESSAGES`
 5803 category may invalidate existing open catalogs.

5804 If a file descriptor is used to implement message catalog descriptors, the `FD_CLOEXEC` flag
 5805 shall be set; see `<fcntl.h>`.

5806 If the value of the `oflag` argument is 0, the `LANG` environment variable is used to locate the
 5807 catalog without regard to the `LC_MESSAGES` category. If the `oflag` argument is
 5808 `NL_CAT_LOCALE`, the `LC_MESSAGES` category is used to locate the message catalog (see the
 5809 Base Definitions volume of IEEE Std. 1003.1-200x, Section 8.2, Internationalization Variables).

5810 **RETURN VALUE**

5811 Upon successful completion, `catopen()` shall return a message catalog descriptor for use on
 5812 subsequent calls to `catgets()` and `catclose()`. Otherwise, `catopen()` shall return `(nl_catd) -1` and set
 5813 `errno` to indicate the error.

5814 **ERRORS**5815 The `catopen()` function may fail if:

5816 [EACCES] Search permission is denied for the component of the path prefix of the
 5817 message catalog or read permission is denied for the message catalog.

5818 [EMFILE] {OPEN_MAX} file descriptors are currently open in the calling process.

5819 [ENAMETOOLONG]

5820 The length of a path name of the message catalog exceeds {PATH_MAX} or a
 5821 path name component is longer than {NAME_MAX}.

5822 [ENAMETOOLONG]

5823 Path name resolution of a symbolic link produced an intermediate result
 5824 whose length exceeds {PATH_MAX}.

5825 [ENFILE] Too many files are currently open in the system.

5826 [ENOENT] The message catalog does not exist or the `name` argument points to an empty
 5827 string.

5828 [ENOMEM] Insufficient storage space is available.

- 5829 [ENOTDIR] A component of the path prefix of the message catalog is not a directory.
- 5830 **EXAMPLES**
- 5831 None.
- 5832 **APPLICATION USAGE**
- 5833 Some implementations of *catopen()* use *malloc()* to allocate space for internal buffer areas. The
- 5834 *catopen()* function may fail if there is insufficient storage space available to accommodate these
- 5835 buffers.
- 5836 Portable applications must assume that message catalog descriptors are not valid after a call to
- 5837 one of the *exec* functions.
- 5838 Application writers should be aware that guidelines for the location of message catalogs have
- 5839 not yet been developed. Therefore they should take care to avoid conflicting with catalogs used
- 5840 by other applications and the standard utilities.
- 5841 **RATIONALE**
- 5842 None.
- 5843 **FUTURE DIRECTIONS**
- 5844 None.
- 5845 **SEE ALSO**
- 5846 *catclose()*, *catgets()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<fcntl.h>`,
- 5847 `<nl_types.h>`, the Shell and Utilities volume of IEEE Std. 1003.1-200x
- 5848 **CHANGE HISTORY**
- 5849 First released in Issue 2.
- 5850 **Issue 4**
- 5851 The type of argument *name* is changed from **char*** to **const char***.
- 5852 The DESCRIPTION is updated:
- 5853
- To indicate the longevity of message catalog descriptors.
 - To specify values for the *oflag* argument and the effect of *LC_MESSAGES* and *NLSPATH*.
- 5854
- 5855 The [EACCES], [EMFILE], [ENAMETOOLONG], [ENFILE], [ENOENT], and [ENOTDIR] errors
- 5856 are added to the ERRORS section.
- 5857 The APPLICATION USAGE section is updated to indicate:
- 5858
- Portable applications should not assume the continued validity of message catalog
 - descriptors after a call to one of the *exec* functions.
- 5859
- 5860
- Message catalogs must be located with care.
- 5861 **Issue 4, Version 2**
- 5862 The following change is incorporated for X/OPEN UNIX conformance:
- 5863
- In the ERRORS section, an [ENAMETOOLONG] condition is defined that may report
 - excessive length of an intermediate result of path name resolution of a symbolic link.
- 5864

5865 **NAME**

5866 cbrt, cbrtf, cbrtl — cube root functions

5867 **SYNOPSIS**

5868 #include <math.h>

5869 double cbrt(double x);

5870 float cbrtf(float x);

5871 long double cbrtl(long double x);

5872 **DESCRIPTION**

5873 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
 5874 conflict between the requirements described here and the ISO C standard is unintentional. This
 5875 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

5876 These functions shall compute the real cube root of x .

5877 An application wishing to check for error situations should set *errno* to 0 before calling these
 5878 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

5879 **RETURN VALUE**5880 Upon successful completion, these functions shall return the cube root of x .5881 If x is $\pm\text{Inf}$, these functions shall return x .5882 If x is NaN, NaN shall be returned and *errno* may be set to [EDOM].5883 **ERRORS**

5884 These functions may fail if:

5885 [EDOM] The value of x is NaN.5886 **EXAMPLES**

5887 None.

5888 **APPLICATION USAGE**

5889 None.

5890 **RATIONALE**

5891 For some applications, a true cube root function, which returns negative results for negative
 5892 arguments, is more appropriate than $\text{pow}(x, 1.0/3.0)$, which returns a NaN for x less than 0.

5893 **FUTURE DIRECTIONS**

5894 None.

5895 **SEE ALSO**

5896 The Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

5897 **CHANGE HISTORY**

5898 First released in Issue 4, Version 2.

5899 **Issue 5**

5900 Moved from X/OPEN UNIX extension to BASE.

5901 **Issue 6**5902 The *cbrt()* function is no longer marked XSI.5903 The *cbrtf()* and *cbrtl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

5904 **NAME**

5905 ccos, ccosf, ccosl — complex cosine functions

5906 **SYNOPSIS**

5907 #include <complex.h>

5908 double complex ccos(double complex *z*);5909 float complex ccosf(float complex *z*);5910 long double complex ccosl(long double complex *z*);5911 **DESCRIPTION**

5912 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5913 conflict between the requirements described here and the ISO C standard is unintentional. This
5914 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

5915 These functions shall compute the complex cosine of *z*.5916 **RETURN VALUE**

5917 These functions shall return the complex cosine value.

5918 **ERRORS**

5919 No errors are defined.

5920 **EXAMPLES**

5921 None.

5922 **APPLICATION USAGE**

5923 None.

5924 **RATIONALE**

5925 None.

5926 **FUTURE DIRECTIONS**

5927 None.

5928 **SEE ALSO**5929 *cacos()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>5930 **CHANGE HISTORY**

5931 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5932 **NAME**

5933 ccosh, ccoshf, ccoshl — complex hyperbolic cosine functions

5934 **SYNOPSIS**

5935 #include <complex.h>

5936 double complex ccosh(double complex *z*);5937 float complex ccoshf(float complex *z*);5938 long double complex ccoshl(long double complex *z*);5939 **DESCRIPTION**

5940 cx The functionality described on this reference page is aligned with the ISO C standard. Any
5941 conflict between the requirements described here and the ISO C standard is unintentional. This
5942 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

5943 These functions shall compute the complex hyperbolic cosine of *z*.5944 **RETURN VALUE**

5945 These functions shall return the complex hyperbolic cosine value.

5946 **ERRORS**

5947 No errors are defined.

5948 **EXAMPLES**

5949 None.

5950 **APPLICATION USAGE**

5951 None.

5952 **RATIONALE**

5953 None.

5954 **FUTURE DIRECTIONS**

5955 None.

5956 **SEE ALSO**5957 *cacosh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>5958 **CHANGE HISTORY**

5959 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

5960 **NAME**

5961 ceil, ceilf, ceill — ceiling value function

5962 **SYNOPSIS**

5963 #include <math.h>

5964 double ceil(double x);

5965 float ceilf(float x);

5966 long double ceill(long double x);

5967 **DESCRIPTION**5968 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
5969 conflict between the requirements described here and the ISO C standard is unintentional. This
5970 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.5971 The *ceil()*, *ceilf()*, and *ceill()* functions shall compute the smallest integral value not less than *x*.5972 An application wishing to check for error situations should set *errno* to 0 before calling *ceil()*. If
5973 *errno* is non-zero on return, or the return value is NaN, an error has occurred.5974 **RETURN VALUE**5975 Upon successful completion, the *ceil()*, *ceilf()*, and *ceill()* functions shall return the smallest
5976 integral value not less than *x*, expressed as a type **double**.5977 **XSI** If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].5978 If the correct value would cause overflow, HUGE_VAL shall be returned and *errno* set to
5979 **XSI** [ERANGE]. If *x* is $\pm\text{Inf}$ or ± 0 , the value of *x* shall be returned.5980 **ERRORS**5981 The *ceil()*, *ceilf()*, and *ceill()* functions shall fail if:

5982 [ERANGE] The result overflows.

5983 The *ceil()*, *ceilf()*, and *ceill()* functions may fail if:5984 **XSI** [EDOM] The value of *x* is NaN.5985 **XSI** No other errors shall occur.5986 **EXAMPLES**

5987 None.

5988 **APPLICATION USAGE**5989 The integral value returned by *ceil()* as a **double** need not be expressible as an **int** or **long**. The
5990 return value should be tested before assigning it to an integer type to avoid the undefined results
5991 of an integer overflow.5992 The *ceil()* function can only overflow when the floating point representation has
5993 DBL_MANT_DIG > DBL_MAX_EXP.5994 **RATIONALE**

5995 None.

5996 **FUTURE DIRECTIONS**

5997 None.

5998 **SEE ALSO**5999 *floor()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

6000 **CHANGE HISTORY**

6001 First released in Issue 1. Derived from Issue 1 of the SVID.

6002 **Issue 4**

6003 References to *matherr()* are removed.

6004 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the
6005 ISO C standard and to rationalize error handling in the mathematics functions.

6006 The return value specified for [EDOM] is marked as an extension.

6007 Support for x being $\pm\text{Inf}$ or ± 0 is added to the RETURN VALUE section and marked as an
6008 extension.

6009 **Issue 5**

6010 The DESCRIPTION is updated to indicate how an application should check for an error. This
6011 text was previously published in the APPLICATION USAGE section.

6012 **Issue 6**

6013 The *ceilf()* and *ceilll()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

6014 NAME

6015 `cexp`, `cexpf`, `cexpl` — complex exponential functions

6016 SYNOPSIS

6017 `#include <complex.h>`

6018 `double complex cexp(double complex z);`

6019 `float complex cexpf(float complex z);`

6020 `long double complex cexpl(long double complex z);`

6021 DESCRIPTION

6022 `CX` The functionality described on this reference page is aligned with the ISO C standard. Any
6023 conflict between the requirements described here and the ISO C standard is unintentional. This
6024 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

6025 These functions shall compute the complex exponent of z , defined as e^z .

6026 RETURN VALUE

6027 These functions shall return the complex exponential value of z .

6028 ERRORS

6029 No errors are defined.

6030 EXAMPLES

6031 None.

6032 APPLICATION USAGE

6033 None.

6034 RATIONALE

6035 None.

6036 FUTURE DIRECTIONS

6037 None.

6038 SEE ALSO

6039 `clog()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<complex.h>`

6040 CHANGE HISTORY

6041 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

6042 **NAME**

6043 cfgetispeed — get input baud rate

6044 **SYNOPSIS**

6045 #include <termios.h>

6046 speed_t cfgetispeed(const struct termios *termios_p);

6047 **DESCRIPTION**6048 The *cfgetispeed()* function shall extract the input baud rate from the **termios** structure to which
6049 the *termios_p* argument points.6050 This function shall return exactly the value in the **termios** data structure, without interpretation.6051 **RETURN VALUE**6052 Upon successful completion, *cfgetispeed()* shall return a value of type **speed_t** representing the
6053 input baud rate.6054 **ERRORS**

6055 No errors are defined.

6056 **EXAMPLES**

6057 None.

6058 **APPLICATION USAGE**

6059 None.

6060 **RATIONALE**6061 The term *baud* is used historically here, but is not technically correct. This is properly “bits per
6062 second”, which may not be the same as baud. However, the term is used because of the
6063 historical usage and understanding.6064 The *cfgetospeed()*, *cfgetispeed()*, *cfsetospeed()*, and *cfsetispeed()* functions do not take arguments as
6065 numbers, but rather as symbolic names. There are two reasons for this:

- 6066 1. Historically, numbers were not used because of the way the rate was stored in the data
-
- 6067 structure. This is retained even though a function is now used.
-
- 6068 2. More importantly, only a limited set of possible rates is at all portable, and this constrains
-
- 6069 the application to that set.

6070 There is nothing to prevent an implementation to accept, as an extension, a number (such as 126)
6071 if it wished, and because the encoding of the Bxxx symbols is not specified, this can be done so
6072 no ambiguity is introduced.6073 Setting the input baud rate to zero was a mechanism to allow for split baud rates. Clarifications
6074 in this volume of IEEE Std. 1003.1-200x have made it possible to determine whether split rates
6075 are supported and to support them without having to treat zero as a special case. Since this
6076 functionality is also confusing, it has been declared obsolescent. The 0 argument referred to is
6077 the literal constant 0, not the symbolic constant B0. This volume of IEEE Std. 1003.1-200x does
6078 not preclude B0 from being defined as the value 0; in fact, implementations would likely benefit
6079 from the two being equivalent. This volume of IEEE Std. 1003.1-200x does not fully specify
6080 whether the previous *cfsetispeed()* value is retained after a *tcgetattr()* as the actual value or as
6081 zero. Therefore, portable applications should always set both the input speed and output speed
6082 when setting either.6083 In historical implementations, the baud rate information is traditionally kept in **c_cflag**.
6084 Applications should be written to presume that this might be the case (and thus not blindly copy
6085 **c_cflag**), but not to rely on it in case it is in some other field of the structure. Setting the **c_cflag**
6086 field absolutely after setting a baud rate is a non-portable action because of this. In general, the

6087 unused parts of the flag fields might be used by the implementation and should not be blindly
6088 copied from the descriptions of one terminal device to another.

6089 **FUTURE DIRECTIONS**

6090 None.

6091 **SEE ALSO**

6092 *cfgetospeed()*, *cfsetispeed()*, *cfsetospeed()*, *tcgetattr()*, the Base Definitions volume of
6093 IEEE Std. 1003.1-200x, `<termios.h>`, the Base Definitions volume of IEEE Std. 1003.1-200x,
6094 Chapter 11, General Terminal Interface

6095 **CHANGE HISTORY**

6096 First released in Issue 3.

6097 Entry included for alignment with the POSIX.1-1988 standard.

6098 **Issue 4**

6099 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 6100
- The type of the argument *termios_p* is changed from **struct termios*** to **const struct termios***.
 - The DESCRIPTION is changed to indicate that the function simply returns the value from *termios_p*, irrespective of how that structure was obtained. Issue 3 states that if *termios_p* was not obtained by a successful call to *tcgetattr()*, the behavior is undefined.
- 6101
6102
6103

6104 **NAME**6105 `cfgetospeed` — get output baud rate6106 **SYNOPSIS**6107 `#include <termios.h>`6108 `speed_t cfgetospeed(const struct termios *termios_p);`6109 **DESCRIPTION**6110 The `cfgetospeed()` function shall extract the output baud rate from the **termios** structure to which
6111 the `termios_p` argument points.6112 This function shall return exactly the value in the **termios** data structure, without interpretation.6113 **RETURN VALUE**6114 Upon successful completion, `cfgetospeed()` shall return a value of type **speed_t** representing the
6115 output baud rate.6116 **ERRORS**

6117 No errors are defined.

6118 **EXAMPLES**

6119 None.

6120 **APPLICATION USAGE**

6121 None.

6122 **RATIONALE**6123 Refer to `cfgetispeed()`.6124 **FUTURE DIRECTIONS**

6125 None.

6126 **SEE ALSO**6127 `cfgetispeed()`, `cfsetispeed()`, `cfsetospeed()`, `tcgetattr()`, the Base Definitions volume of
6128 IEEE Std. 1003.1-200x, `<termios.h>`, the Base Definitions volume of IEEE Std. 1003.1-200x,
6129 Chapter 11, General Terminal Interface6130 **CHANGE HISTORY**

6131 First released in Issue 3.

6132 Entry included for alignment with the POSIX.1-1988 standard.

6133 **Issue 4**

6134 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 6135
- The type of the argument `termios_p` is changed from **struct termios*** to **const struct termios***.
 - The DESCRIPTION is changed to indicate that the function simply returns the value from `termios_p`, irrespective of how that structure was obtained. Issue 3 states that if `termios_p` was not obtained by a successful call to `tcgetattr()`, the behavior is undefined.
- 6136
-
- 6137
-
- 6138

6139 **NAME**

6140 cfsetispeed — set input baud rate

6141 **SYNOPSIS**

6142 #include <termios.h>

6143 int cfsetispeed(struct termios *termios_p, speed_t speed);

6144 **DESCRIPTION**

6145 The *cfsetispeed()* function shall set the input baud rate stored in the structure pointed to by
6146 *termios_p* to *speed*.

6147 There shall be no effect on the baud rates set in the hardware until a subsequent successful call
6148 to *tcsetattr()* with the same **termios** structure. Similarly, errors resulting from attempts to set
6149 baud rates not supported by the terminal device need not be detected until the *tcsetattr()*
6150 function is called.

6151 **RETURN VALUE**

6152 Upon successful completion, *cfsetispeed()* shall return 0; otherwise, -1 shall be returned, and
6153 *errno* may be set to indicate the error.

6154 **ERRORS**

6155 The *cfsetispeed()* function may fail if:

6156 [EINVAL] The *speed* value is not a valid baud rate.

6157 [EINVAL] The value of *speed* is outside the range of possible speed values as specified in
6158 <**termios.h**>.

6159 **EXAMPLES**

6160 None.

6161 **APPLICATION USAGE**

6162 None.

6163 **RATIONALE**

6164 Refer to *cfgetispeed()*.

6165 **FUTURE DIRECTIONS**

6166 None.

6167 **SEE ALSO**

6168 *cfgetispeed()*, *cfgetospeed()*, *cfsetospeed()*, *tcsetattr()*, the Base Definitions volume of
6169 IEEE Std. 1003.1-200x, <**termios.h**>, the Base Definitions volume of IEEE Std. 1003.1-200x,
6170 Chapter 11, General Terminal Interface

6171 **CHANGE HISTORY**

6172 First released in Issue 3.

6173 Entry included for alignment with the POSIX.1-1988 standard.

6174 **Issue 4**

6175 The first description of the [EINVAL] error is added and is marked as an extension.

6176 **Issue 4, Version 2**

6177 The ERRORS section is changed to indicate that [EINVAL] may be returned if the specified
6178 speed is outside the range of possible speed values given in <**termios.h**>.

6179 **Issue 6**

6180 The following new requirements on POSIX implementations derive from alignment with the
6181 Single UNIX Specification:

- 6182
- The optional setting of *errno* and the [EINVAL] error conditions are added.

6183 **NAME**

6184 cfsetospeed — set output baud rate

6185 **SYNOPSIS**

6186 #include <termios.h>

6187 int cfsetospeed(struct termios *termios_p, speed_t speed);

6188 **DESCRIPTION**6189 The *cfsetospeed()* function shall set the output baud rate stored in the structure pointed to by
6190 *termios_p* to *speed*.6191 There shall be no effect on the baud rates set in the hardware until a subsequent successful call
6192 to *tcsetattr()* with the same **termios** structure. Similarly, errors resulting from attempts to set
6193 baud rates not supported by the terminal device need not be detected until the *tcsetattr()*
6194 function is called.6195 **RETURN VALUE**6196 Upon successful completion, *cfsetospeed()* shall return 0; otherwise, it shall return -1 and *errno*
6197 may be set to indicate the error.6198 **ERRORS**6199 The *cfsetospeed()* function may fail if:6200 [EINVAL] The *speed* value is not a valid baud rate.6201 [EINVAL] The value of *speed* is outside the range of possible speed values as specified in
6202 <termios.h>.6203 **EXAMPLES**

6204 None.

6205 **APPLICATION USAGE**

6206 None.

6207 **RATIONALE**6208 Refer to *cfgetispeed()*.6209 **FUTURE DIRECTIONS**

6210 None.

6211 **SEE ALSO**6212 *cfgetispeed()*, *cfgetospeed()*, *cfsetispeed()*, *tcsetattr()*, the Base Definitions volume of
6213 IEEE Std. 1003.1-200x, <termios.h>, the Base Definitions volume of IEEE Std. 1003.1-200x,
6214 Chapter 11, General Terminal Interface6215 **CHANGE HISTORY**

6216 First released in Issue 3.

6217 Entry included for alignment with the POSIX.1-1988 standard.

6218 **Issue 4**

6219 The first description of the [EINVAL] error is added and is marked as an extension.

6220 **Issue 4, Version 2**6221 The ERRORS section is changed to indicate that [EINVAL] may be returned if the specified
6222 speed is outside the range of possible speed values given in <termios.h>.

6223 **Issue 6**

6224 The following new requirements on POSIX implementations derive from alignment with the
6225 Single UNIX Specification:

- 6226 • The optional setting of *errno* and the [EINVAL] error conditions are added.

6227 **NAME**

6228 chdir — change working directory

6229 **SYNOPSIS**

6230 #include <unistd.h>

6231 int chdir(const char *path);

6232 **DESCRIPTION**

6233 The *chdir()* function shall cause the directory named by the path name pointed to by the *path*
 6234 argument to become the current working directory; that is, the starting point for path searches
 6235 for path names not beginning with '/ '.

6236 **RETURN VALUE**

6237 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned, the current
 6238 working directory shall remain unchanged, and *errno* shall be set to indicate the error.

6239 **ERRORS**6240 The *chdir()* function shall fail if:

6241 [EACCES] Search permission is denied for any component of the path name.

6242 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*
6243 argument.

6244 [ENAMETOOLONG]

6245 The length of the *path* argument exceeds {PATH_MAX} or a path name
6246 component is longer than {NAME_MAX}.6247 [ENOENT] A component of *path* does not name an existing directory or *path* is an empty
6248 string.

6249 [ENOTDIR] A component of the path name is not a directory.

6250 The *chdir()* function may fail if:6251 [ELOOP] More than {SYMLOOP_MAX} symbolic links were encountered during
6252 resolution of the *path* argument.

6253 [ENAMETOOLONG]

6254 As a result of encountering a symbolic link in resolution of the *path* argument,
6255 the length of the substituted path name string exceeded {PATH_MAX}.6256 **EXAMPLES**6257 **Changing the Current Working Directory**6258 The following example makes the value pointed to by **directory**, **/tmp**, the current working
6259 directory.

6260 #include <unistd.h>

6261 ...

6262 char *directory = "/tmp";

6263 int ret;

6264 ret = chdir (directory);

6265 **APPLICATION USAGE**

6266 The *chdir()* function only affects the working directory of the current process. The result if a
 6267 NULL argument is passed to *chdir()* is unspecified because some implementations dynamically
 6268 allocate space in that case.

6269 **RATIONALE**

6270 The *chdir()* function only affects the working directory of the current process.

6271 The result if a NULL argument is passed to *chdir()* is left implementation-defined because some
 6272 implementations dynamically allocate space in that case.

6273 **FUTURE DIRECTIONS**

6274 None.

6275 **SEE ALSO**

6276 *getcwd()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

6277 **CHANGE HISTORY**

6278 First released in Issue 1. Derived from Issue 1 of the SVID.

6279 **Issue 4**

6280 The <**unistd.h**> header is added to the SYNOPSIS section.

6281 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 6282 • The type of argument *path* is changed from **char*** to **const char***.

6283 The following change is incorporated for alignment with the FIPS requirements:

- 6284 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path
 6285 name component is larger than {NAME_MAX} is now defined as mandatory and marked as
 6286 an extension.

6287 **Issue 4, Version 2**

6288 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 6289 • It states that [ELOOP] is returned if too many symbolic links are encountered during path
 6290 name resolution.
- 6291 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an
 6292 intermediate result of path name resolution of a symbolic link.

6293 **Issue 6**

6294 The APPLICATION USAGE section is added.

6295 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 6296 • The [ENAMETOOLONG] error is restored as an error dependent on **_POSIX_NO_TRUNC**.
 6297 This is since behavior may vary from one file system to another.

6298 The following new requirements on POSIX implementations derive from alignment with the
 6299 Single UNIX Specification:

- 6300 • The [ELOOP] mandatory error condition is added.
- 6301 • A second [ENAMETOOLONG] is added as an optional error condition.

6302 The following changes were made to align with the IEEE P1003.1a draft standard:

- 6303 • The [ELOOP] optional error condition is added.

6304 **NAME**

6305 chmod — change mode of a file

6306 **SYNOPSIS**

6307 #include <sys/stat.h>

6308 int chmod(const char *path, mode_t mode);

6309 **DESCRIPTION**

6310 XSI The *chmod()* function shall change S_ISUID, S_ISGID, S_ISVTX, and the file permission bits of
6311 the file named by the path name pointed to by the *path* argument to the corresponding bits in the
6312 *mode* argument. The application shall ensure that the effective user ID of the process matches the
6313 owner of the file or the process has appropriate privileges in order to do this.

6314 S_ISUID, S_ISGID, and the file permission bits are described in <sys/stat.h>.

6315 XSI If a directory is writable and the mode bit S_ISVTX is set on the directory, a process may remove
6316 or rename files within that directory only if one or more of the following is true:

- 6317 • The effective user ID of the process is the same as that of the owner ID of the file.
- 6318 • The effective user ID of the process is the same as that of the owner ID of the directory.
- 6319 • The process has appropriate privileges.

6321 XSI If the S_ISVTX bit is set on a non-directory file, the behavior is unspecified.

6322 If the calling process does not have appropriate privileges, and if the group ID of the file does
6323 not match the effective group ID or one of the supplementary group IDs and if the file is a
6324 regular file, bit S_ISGID (set-group-ID on execution) in the file's mode shall be cleared upon
6325 successful return from *chmod()*.

6326 Additional implementation-defined restrictions may cause the S_ISUID and S_ISGID bits in
6327 *mode* to be ignored.

6328 The effect on file descriptors for files open at the time of a call to *chmod()* is implementation-
6329 defined.

6330 Upon successful completion, *chmod()* shall mark for update the *st_ctime* field of the file.

6331 **RETURN VALUE**

6332 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to
6333 indicate the error. If -1 is returned, no change to the file mode occurs.

6334 **ERRORS**

6335 The *chmod()* function shall fail if:

- 6336 [EACCES] Search permission is denied on a component of the path prefix.
- 6337 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*
6338 argument.
- 6339 [ENAMETOOLONG]
6340 The length of the *path* argument exceeds {PATH_MAX} or a path name
6341 component is longer than {NAME_MAX}.
- 6342 [ENOTDIR] A component of the path prefix is not a directory.
- 6343 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.
- 6344 [EPERM] The effective user ID does not match the owner of the file and the process
6345 does not have appropriate privileges.

6346	[EROFS]	The named file resides on a read-only file system.
6347		The <i>chmod()</i> function may fail if:
6348	[EINTR]	A signal was caught during execution of the function.
6349	[EINVAL]	The value of the <i>mode</i> argument is invalid.
6350	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during
6351		resolution of the <i>path</i> argument.
6352	[ENAMETOOLONG]	
6353		As a result of encountering a symbolic link in resolution of the <i>path</i> argument,
6354		the length of the substituted path name strings exceeded {PATH_MAX}.

6355 EXAMPLES

6356 Setting Read Permissions for User, Group, and Others

6357 The following example sets read permissions for the owner, group, and others.

```
6358 #include <sys/stat.h>
6359 const char *path;
6360 ...
6361 chmod(path, S_IRUSR | S_IRGRP | S_IROTH);
```

6362 Setting Read, Write, and Execute Permissions for the Owner Only

6363 The following example sets read, write, and execute permissions for the owner, and no
6364 permissions for group and others.

```
6365 #include <sys/stat.h>
6366 const char *path;
6367 ...
6368 chmod(path, S_IRWXU);
```

6369 Setting Different Permissions for Owner, Group, and Other

6370 The following example sets owner permissions for *CHANGEFILE* to read, write, and execute,
6371 group permissions to read and execute, and other permissions to read.

```
6372 #include <sys/stat.h>
6373 #define CHANGEFILE "/etc/myfile"
6374 ...
6375 chmod(CHANGEFILE, S_IRWXU | S_IRGRP | S_IXGRP | S_IROTH);
```

6376 Setting and Checking File Permissions

6377 The following example sets the file permission bits for a file named */home/cnd/mod1*, then calls
6378 the *stat()* function to verify the permissions.

```
6379 #include <sys/types.h>
6380 #include <sys/stat.h>
6381 int status;
6382 struct stat buffer
6383 ...
```

```
6384     chmod("home/cnd/mod1", S_IRWXU|S_IRWXG|S_IROTH|S_IWOTH);
6385     status = stat("home/cnd/mod1", &buffer);
```

6386 APPLICATION USAGE

6387 In order to ensure that the S_ISUID and S_ISGID bits are set, an application requiring this should
6388 use *stat()* after a successful *chmod()* to verify this.

6389 Any file descriptors currently open by any process on the file could possibly become invalid if
6390 the mode of the file is changed to a value which would deny access to that process. One
6391 situation where this could occur is on a stateless file system. This behavior will not occur in a
6392 conforming environment.

6393 RATIONALE

6394 This volume of IEEE Std. 1003.1-200x specifies that the S_ISGID bit is cleared by *chmod()* on a
6395 regular file under certain conditions. This is specified on the assumption that regular files may
6396 be executed, and the system should prevent users from making executable *setgid()* files perform
6397 with privileges that the caller does not have. On implementations that support execution of
6398 other file types, the S_ISGID bit should be cleared for those file types under the same
6399 circumstances.

6400 Implementations that use the S_ISUID bit to indicate some other function (for example,
6401 mandatory record locking) on non-executable files need not clear this bit on writing. They
6402 should clear the bit for executable files and any other cases where the bit grants special powers
6403 to processes that change the file contents. Similar comments apply to the S_ISGID bit.

6404 FUTURE DIRECTIONS

6405 None.

6406 SEE ALSO

6407 *chown()*, *mkdir()*, *mkfifo()*, *open()*, *stat()*, *statvfs()*, the Base Definitions volume of
6408 IEEE Std. 1003.1-200x, <sys/stat.h>, <sys/types.h>

6409 CHANGE HISTORY

6410 First released in Issue 1. Derived from Issue 1 of the SVID.

6411 Issue 4

6412 The <sys/types.h> header is now marked as optional (OH); this header need not be included on
6413 XSI-conformant systems.

6414 The [EINVAL] error is marked as an extension.

6415 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 6416 • The type of argument *path* is changed from **char*** to **const char***.

6417 The following change is incorporated for alignment with the FIPS requirements:

- 6418 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path
6419 name component is larger than {NAME_MAX} is now defined as mandatory and marked as
6420 an extension.

6421 Issue 4, Version 2

6422 The following changes are incorporated for X/OPEN UNIX conformance:

- 6423 • The DESCRIPTION is updated to describe X/OPEN UNIX functionality relating to
6424 permission checks applied when removing or renaming files in a directory having the
6425 S_ISVTX bit set.
- 6426 • In the ERRORS section, the condition whereby [ELOOP] is returned if too many symbolic
6427 links are encountered during path name resolution is defined as mandatory, and [EINTR] is

- 6428 added as an optional error.
- 6429 • In the ERRORS section, a second [ENAMETOOLONG] condition is defined that may report
6430 excessive length of an intermediate result of path name resolution of a symbolic link.
- 6431 **Issue 6**
- 6432 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:
- 6433 • The [ENAMETOOLONG] error is restored as an error dependent on _POSIX_NO_TRUNC.
6434 This is since behavior may vary from one file system to another.
- 6435 The following new requirements on POSIX implementations derive from alignment with the
6436 Single UNIX Specification:
- 6437 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was
6438 required for conforming implementations of previous POSIX specifications, it was not
6439 required for UNIX applications.
- 6440 • The [EINVAL] and [EINTR] optional error conditions are added.
- 6441 • A second [ENAMETOOLONG] is added as an optional error condition.
- 6442 The following changes were made to align with the IEEE P1003.1a draft standard:
- 6443 • The [ELOOP] optional error condition is added.
- 6444 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

6445 **NAME**

6446 chown — change owner and group of a file

6447 **SYNOPSIS**

6448 #include <unistd.h>

6449 int chown(const char *path, uid_t owner, gid_t group);

6450 **DESCRIPTION**6451 The *path* argument points to a path name naming a file. The user ID and group ID of the named
6452 file are set to the numeric values contained in *owner* and *group*, respectively.6453 Only processes with an effective user ID equal to the user ID of the file or with appropriate
6454 privileges may change the ownership of a file. If `_POSIX_CHOWN_RESTRICTED` is in effect for
6455 *path*:

- 6456
- Changing the user ID is restricted to processes with appropriate privileges.
 - Changing the group ID is permitted to a process with an effective user ID equal to the user
6457 ID of the file, but without appropriate privileges, if and only if *owner* is equal to the file's user
6458 ID or `(uid_t)-1` and *group* is equal either to the calling process' effective group ID or to one of
6459 its supplementary group IDs.

6461 If the specified file is a regular file, one or more of the `S_IXUSR`, `S_IXGRP`, or `S_IXOTH` bits of
6462 the file mode are set, and the process does not have appropriate privileges, the set-user-ID
6463 (`S_ISUID`) and set-group-ID (`S_ISGID`) bits of the file mode shall be cleared upon successful
6464 return from *chown*(*path*). If the specified file is a regular file, one or more of the `S_IXUSR`, `S_IXGRP`,
6465 or `S_IXOTH` bits of the file mode are set, and the process has appropriate privileges, it is
6466 implementation-defined whether the set-user-ID and set-group-ID bits are altered. If the *chown*(*path*)
6467 function is successfully invoked on a file that is not a regular file and one or more of the
6468 `S_IXUSR`, `S_IXGRP`, or `S_IXOTH` bits of the file mode are set, the set-user-ID and set-group-ID
6469 bits may be cleared.6470 If *owner* or *group* is specified as `(uid_t)-1` or `(gid_t)-1`, respectively, the corresponding ID of the
6471 file is unchanged.6472 Upon successful completion, *chown*(*path*) shall mark for update the *st_ctime* field of the file.6473 **RETURN VALUE**6474 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to
6475 indicate the error. If -1 is returned, no changes are made in the user ID and group ID of the file.6476 **ERRORS**6477 The *chown*(*path*) function shall fail if:

- 6478 [EACCES] Search permission is denied on a component of the path prefix.
-
- 6479 [ELOOP] A loop exists in symbolic links encountered during resolution of the
- path*
-
- 6480 argument.
-
- 6481 [ENAMETOOLONG] The length of the
- path*
- argument exceeds {PATH_MAX} or a path name
-
- 6482 component is longer than {NAME_MAX}.
-
- 6483 [ENOTDIR] A component of the path prefix is not a directory.
-
- 6484 [ENOENT] A component of
- path*
- does not name an existing file or
- path*
- is an empty string.
-
- 6485 [EPERM] The effective user ID does not match the owner of the file, or the calling
-
- 6486 process does not have appropriate privileges and
-
- 6487
- `_POSIX_CHOWN_RESTRICTED`
- indicates that such privilege is required.
-
- 6488

6489	[EROFS]	The named file resides on a read-only file system.
6490		The <i>chown()</i> function may fail if:
6491	[EIO]	An I/O error occurred while reading or writing to the file system.
6492	[EINTR]	The <i>chown()</i> function was interrupted by a signal which was caught.
6493	[EINVAL]	The owner or group ID supplied is not a value supported by the
6494		implementation.
6495	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during
6496		resolution of the <i>path</i> argument.
6497	[ENAMETOOLONG]	
6498		As a result of encountering a symbolic link in resolution of the <i>path</i> argument,
6499		the length of the substituted path name string exceeded {PATH_MAX}.

6500 EXAMPLES

6501 None.

6502 APPLICATION USAGE

6503 Although *chown()* can be used on some systems by the file owner to change the owner and
 6504 group to any desired values, the only portable use of this function is to change the group of a file
 6505 to the effective GID of the calling process or to a member of its group set.

6506 RATIONALE

6507 System III and System V allow a user to give away files; that is, the owner of a file may change
 6508 its user ID to anything. This is a serious problem for implementations that are intended to meet
 6509 government security regulations. Version 7 and 4.3 BSD permit only the superuser to change the
 6510 user ID of a file. Some government agencies (usually not ones concerned directly with security)
 6511 find this limitation too confining. This volume of IEEE Std. 1003.1-200x uses *may* to permit
 6512 secure implementations while not disallowing System V.

6513 System III and System V allow the owner of a file to change the group ID to anything. Version 7
 6514 permits only the superuser to change the group ID of a file. 4.3 BSD permits the owner to
 6515 change the group ID of a file to its effective group ID or to any of the groups in the list of
 6516 supplementary group IDs, but to no others.

6517 The POSIX.1-1990 standard requires that the *chown()* function invoked by a non-appropriate
 6518 privileged process clear the *S_ISGID* and the *S_ISUID* bits for regular files, and permits them to
 6519 be cleared for other types of files. This is so that changes in accessibility do not accidentally
 6520 cause files to become security holes. Unfortunately, requiring these bits to be cleared on non-
 6521 executable data files also clears the mandatory file locking bit (shared withwith *S_ISGID*), which
 6522 is an extension on many implementations (it first appeared in System V). These bits should only
 6523 be required to be cleared on regular files that have one or more of their execute bits set.

6524 FUTURE DIRECTIONS

6525 None.

6526 SEE ALSO

6527 *chmod()*, *pathconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<sys/types.h>`,
 6528 `<unistd.h>`

6529 CHANGE HISTORY

6530 First released in Issue 1. Derived from Issue 1 of the SVID.

6531 Issue 4

6532 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on
6533 XSI-conformant systems.

6534 The value for *owner* of `(uid_t)-1` is added to the DESCRIPTION to allow the use of `-1` by the
6535 owner of a file to change the group ID only.

6536 The APPLICATION USAGE section is added.

6537 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 6538 • The type of argument *path* is changed from `char*` to `const char*`.

6539 The following changes are incorporated for alignment with the FIPS requirements:

6540 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path
6541 name component is larger than {NAME_MAX} is now defined as mandatory and marked as
6542 an extension.

6543 • In the ERRORS section, the condition whereby [EPERM] is returned when an attempt is
6544 made to change the user ID of a file and the caller does not have appropriate privileges is
6545 now defined as mandatory and marked as an extension.

6546 Issue 4, Version 2

6547 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

6548 • It states that [ELOOP] is returned if too many symbolic links are encountered during path
6549 name resolution.

6550 • The [EIO] and [EINTR] optional conditions are added.

6551 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an
6552 intermediate result of path name resolution of a symbolic link.

6553 Issue 6

6554 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

6555 • The wording describing the optional dependency on `_POSIX_CHOWN_RESTRICTED` is
6556 restored.

6557 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.
6558 This is since behavior may vary from one file system to another.

6559 • The [EPERM] error is restored as an error dependent on `_POSIX_CHOWN_RESTRICTED`.
6560 This is since its operand is a path name and applications should be aware that the error may
6561 not occur for that path name if the file system does not support
6562 `_POSIX_CHOWN_RESTRICTED`.

6563 The following new requirements on POSIX implementations derive from alignment with the
6564 Single UNIX Specification:

6565 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was
6566 required for conforming implementations of previous POSIX specifications, it was not
6567 required for UNIX applications.

6568 • The value for *owner* of `(uid_t)-1` allows the use of `-1` by the owner of a file to change the
6569 group ID only.

6570 • The [ELOOP] mandatory error condition is added.

6571 • The [EIO] and [EINTR] optional error conditions are added.

6572

- A second [ENAMETOOLONG] is added as an optional error condition.

6573

The following changes were made to align with the IEEE P1003.1a draft standard:

6574

- Clarification is added that the S_ISUID and S_ISGID bits do not need to be cleared when the process has appropriate privileges.

6575

6576

- The [ELOOP] optional error condition is added.

6577 **NAME**

6578 cimag, cimagf, cimagl — complex imaginary functions

6579 **SYNOPSIS**

6580 #include <complex.h>

6581 double cimag(double complex z);

6582 float cimagf(float complex z);

6583 long double cimagl(long double complex z);

6584 **DESCRIPTION**

6585 cx The functionality described on this reference page is aligned with the ISO C standard. Any
6586 conflict between the requirements described here and the ISO C standard is unintentional. This
6587 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

6588 These functions shall compute the imaginary part of *z*.6589 **RETURN VALUE**

6590 These functions shall return the imaginary part value (as a real).

6591 **ERRORS**

6592 No errors are defined.

6593 **EXAMPLES**

6594 None.

6595 **APPLICATION USAGE**6596 For a variable *z* of complex type:6597 `z == creal(z) + cimag(z)*I`6598 **RATIONALE**

6599 None.

6600 **FUTURE DIRECTIONS**

6601 None.

6602 **SEE ALSO**6603 *carg()*, *conj()*, *cproj()*, *creal()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>6604 **CHANGE HISTORY**

6605 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

6606 **NAME**

6607 clearerr — clear indicators on a stream

6608 **SYNOPSIS**

6609 #include <stdio.h>

6610 void clearerr(FILE **stream*);6611 **DESCRIPTION**

6612 cx The functionality described on this reference page is aligned with the ISO C standard. Any
6613 conflict between the requirements described here and the ISO C standard is unintentional. This
6614 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

6615 The *clearerr()* function shall clear the end-of-file and error indicators for the stream to which
6616 *stream* points.

6617 **RETURN VALUE**6618 The *clearerr()* function shall return no value.6619 **ERRORS**

6620 No errors are defined.

6621 **EXAMPLES**

6622 None.

6623 **APPLICATION USAGE**

6624 None.

6625 **RATIONALE**

6626 None.

6627 **FUTURE DIRECTIONS**

6628 None.

6629 **SEE ALSO**

6630 The Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>

6631 **CHANGE HISTORY**

6632 First released in Issue 1. Derived from Issue 1 of the SVID.

6633 **NAME**

6634 clock — report CPU time used

6635 **SYNOPSIS**

6636 #include <time.h>

6637 clock_t clock(void);

6638 **DESCRIPTION**

6639 cx The functionality described on this reference page is aligned with the ISO C standard. Any
6640 conflict between the requirements described here and the ISO C standard is unintentional. This
6641 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

6642 The *clock()* function shall return the implementation's best approximation to the processor time
6643 used by the process since the beginning of an implementation-defined time related only to the
6644 process invocation.

6645 **RETURN VALUE**

6646 To determine the time in seconds, the value returned by *clock()* should be divided by the value
6647 xsi of the macro `CLOCKS_PER_SEC`. `CLOCKS_PER_SEC` is defined to be one million in `<time.h>`.
6648 If the processor time used is not available or its value cannot be represented, the function shall
6649 return the value `(clock_t)-1`.

6650 **ERRORS**

6651 No errors are defined.

6652 **EXAMPLES**

6653 None.

6654 **APPLICATION USAGE**

6655 In order to measure the time spent in a program, *clock()* should be called at the start of the
6656 program and its return value subtracted from the value returned by subsequent calls. The value
6657 returned by *clock()* is defined for compatibility across systems that have clocks with different
6658 resolutions. The resolution on any particular system need not be to microsecond accuracy.

6659 The value returned by *clock()* may wrap around on some systems. For example, on a machine
6660 with 32-bit values for `clock_t`, it wraps after 2 147 seconds or 36 minutes.

6661 **RATIONALE**

6662 None.

6663 **FUTURE DIRECTIONS**

6664 None.

6665 **SEE ALSO**

6666 *asctime()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strftime()*, *strptime()*, *time()*, *utime()*,
6667 the Base Definitions volume of IEEE Std. 1003.1-200x, `<time.h>`

6668 **CHANGE HISTORY**

6669 First released in Issue 1. Derived from Issue 1 of the SVID.

6670 **Issue 4**6671 Reference to the resolution of `CLOCKS_PER_SEC` is marked as an extension.6672 The **ERRORS** section is added.

6673 Advice on how to calculate the time spent in a program is added to the **APPLICATION USAGE**
6674 section.

6675 The following changes are incorporated for alignment with the ISO C standard:

- 6676 • The `<time.h>` header is added to the SYNOPSIS section.
- 6677 • The DESCRIPTION and RETURN VALUE sections, though functionally equivalent to Issue
6678 3, are rewritten for clarity and consistency with the ISO C standard. This issue also defines
6679 under what circumstances `(clock_t)-1` can be returned by the function.
- 6680 • The function is no longer marked as an extension.

6681 **NAME**6682 clock_getcpuclockid — access a process CPU-time clock (**REALTIME**)6683 **SYNOPSIS**

6684 CPT #include <time.h>

6685 int clock_getcpuclockid(pid_t pid, clockid_t *clock_id);

6686

6687 **DESCRIPTION**6688 The *clock_getcpuclockid()* function shall return the clock ID of the CPU-time clock of the process
6689 specified by *pid*. If the process described by *pid* exists and the calling process has permission,
6690 the clock ID of this clock shall be returned in *clock_id*.6691 If *pid* is zero, the *clock_getcpuclockid()* function shall return the clock ID of the CPU-time clock of
6692 the process making the call, in *clock_id*.6693 The conditions under which one process has permission to obtain the CPU-time clock ID of
6694 other processes are implementation-defined.6695 **RETURN VALUE**6696 Upon successful completion, *clock_getcpuclockid()* shall return zero; otherwise, an error number
6697 shall be returned to indicate the error.6698 **ERRORS**6699 The *clock_getcpuclockid()* function shall fail if:6700 [EPERM] The requesting process does not have permission to access the CPU-time
6701 clock for the process.6702 The *clock_getcpuclockid()* function may fail if:6703 [ESRCH] No process can be found corresponding to the process specified by *pid*.6704 **EXAMPLES**

6705 None.

6706 **APPLICATION USAGE**6707 The *clock_getcpuclockid()* function is part of the Process CPU-Time Clocks option and need not
6708 be provided on all implementations.6709 **RATIONALE**

6710 None.

6711 **FUTURE DIRECTIONS**

6712 None.

6713 **SEE ALSO**6714 *clock_getres()*, *timer_create()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <time.h>6715 **CHANGE HISTORY**

6716 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

6717 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

6718 NAME

6719 clock_getres, clock_gettime, clock_settime — clock and timer functions (**REALTIME**)

6720 SYNOPSIS

6721 TMR #include <time.h>

```
6722 int clock_getres(clockid_t clock_id, struct timespec *res);
6723 int clock_settime(clockid_t clock_id, const struct timespec *tp);
6724 int clock_gettime(clockid_t clock_id, struct timespec *tp);
6725
```

6726 DESCRIPTION

6727 The resolution of any clock can be obtained by calling *clock_getres()*. Clock resolutions are
 6728 implementation-defined and cannot be set by a process. If the argument *res* is not NULL, the
 6729 resolution of the specified clock shall be stored in the location pointed to by *res*. If *res* is NULL,
 6730 the clock resolution is not returned. If the *time* argument of *clock_settime()* is not a multiple of *res*,
 6731 then the value is truncated to a multiple of *res*.

6732 The *clock_gettime()* function shall return the current value *tp* for the specified clock, *clock_id*.

6733 The *clock_settime()* function shall set the specified clock, *clock_id*, to the value specified by *tp*.
 6734 Time values that are between two consecutive non-negative integer multiples of the resolution
 6735 of the specified clock are truncated down to the smaller multiple of the resolution.

6736 A clock may be system-wide (that is, visible to all processes) or per-process (measuring time that
 6737 is meaningful only within a process). All implementations shall support a *clock_id* of
 6738 CLOCK_REALTIME defined in <time.h>. This clock represents the realtime clock for the
 6739 system. For this clock, the values returned by *clock_gettime()* and specified by *clock_settime()*
 6740 represent the amount of time (in seconds and nanoseconds) since the Epoch. An implementation
 6741 may also support additional clocks. The interpretation of time values for these clocks is
 6742 unspecified.

6743 If the value of the CLOCK_REALTIME clock is set via *clock_settime()*, the new value of the clock
 6744 shall be used to determine the time of expiration for absolute time services based upon the
 6745 CLOCK_REALTIME clock. This applies to the time at which armed absolute timers expire. If the
 6746 absolute time requested at the invocation of such a time service is before the new value of the
 6747 clock, the time service shall expire immediately as if the clock had reached the requested time
 6748 normally.

6749 Setting the value of the CLOCK_REALTIME clock via *clock_settime()* shall have no effect on
 6750 threads that are blocked waiting for a relative time service based upon this clock, including the
 6751 *nanosleep()* function; nor on the expiration of relative timers based upon this clock.
 6752 Consequently, these time services shall expire when the requested relative interval elapses,
 6753 independently of the new or old value of the clock.

6754 MON If the Monotonic Clock option is supported, all implementations shall support a *clock_id* of
 6755 CLOCK_MONOTONIC defined in <time.h>. This clock represents the monotonic clock for the
 6756 system. For this clock, the value returned by *clock_gettime()* represents the amount of time (in
 6757 seconds and nanoseconds) since an unspecified point in the past (for example, system start-up
 6758 time, or the Epoch). This point does not change after system start-up time. The value of the
 6759 CLOCK_MONOTONIC clock cannot be set via *clock_settime()*. This function shall fail if it is
 6760 invoked with a *clock_id* argument of CLOCK_MONOTONIC.

6761 The effect of setting a clock via *clock_settime()* on armed per-process timers associated with a
 6762 clock other than CLOCK_REALTIME is implementation-defined.

6763 CS If the value of the CLOCK_REALTIME clock is set via *clock_settime()*, the new value of the clock
 6764 shall be used to determine the time at which the system shall awaken a thread blocked on an

- 6765 absolute *clock_nanosleep()* call based upon the CLOCK_REALTIME clock. If the absolute time
 6766 requested at the invocation of such a time service is before the new value of the clock, the call
 6767 shall return immediately as if the clock had reached the requested time normally.
- 6768 Setting the value of the CLOCK_REALTIME clock via *clock_settime()* shall have no effect on any
 6769 thread that is blocked on a relative *clock_nanosleep()* call. Consequently, the call shall return
 6770 when the requested relative interval elapses, independently of the new or old value of the clock.
- 6771 The appropriate privilege to set a particular clock is implementation-defined.
- 6772 CPT If `_POSIX_CPUTIME` is defined, implementations shall support clock ID values obtained by
 6773 invoking *clock_getcpuclockid()*, which represent the CPU-time clock of a given process.
 6774 Implementations shall also support the special `clockid_t` value
 6775 `CLOCK_PROCESS_CPUTIME_ID`, which represents the CPU-time clock of the calling process
 6776 when invoking one of the *clock_**() or *timer_**() functions. For these clock IDs, the values
 6777 returned by *clock_gettime()* and specified by *clock_settime()* represent the amount of execution
 6778 time of the process associated with the clock. Changing the value of a CPU-time clock via
 6779 *clock_settime()* shall have no effect on the behavior of the sporadic server scheduling policy (see
 6780 **Scheduling Policies** (on page 546)).
- 6781 TCT If `_POSIX_THREAD_CPUTIME` is defined, implementations shall support clock ID values
 6782 obtained by invoking *pthread_getcpuclockid()*, which represent the CPU-time clock of a given
 6783 thread. Implementations shall also support the special `clockid_t` value
 6784 `CLOCK_THREAD_CPUTIME_ID`, which represents the CPU-time clock of the calling thread
 6785 when invoking one of the *clock_**() or *timer_**() functions. For these clock IDs, the values
 6786 returned by *clock_gettime()* and specified by *clock_settime()* represent the amount of execution
 6787 time of the thread associated with the clock. Changing the value of a CPU-time clock via
 6788 *clock_settime()* shall have no effect on the behavior of the sporadic server scheduling policy (see
 6789 **Scheduling Policies** (on page 546)).
- 6790 **RETURN VALUE**
- 6791 A return value of 0 shall indicate that the call succeeded. A return value of -1 shall indicate that
 6792 an error occurred, and *errno* shall be set to indicate the error.
- 6793 **ERRORS**
- 6794 The *clock_getres()*, *clock_gettime()*, and *clock_settime()* functions shall fail if:
- 6795 [EINVAL] The *clock_id* argument does not specify a known clock.
- 6796 The *clock_settime()* function shall fail if:
- 6797 [EINVAL] The *tp* argument to *clock_settime()* is outside the range for the given clock ID.
- 6798 [EINVAL] The *tp* argument specified a nanosecond value less than zero or greater than
 6799 or equal to 1 000 million.
- 6800 MON [EINVAL] The value of the *clock_id* argument is `CLOCK_MONOTONIC`.
- 6801 The *clock_settime()* function may fail if:
- 6802 [EPERM] The requesting process does not have the appropriate privilege to set the
 6803 specified clock.

6804 **EXAMPLES**

6805 None.

6806 **APPLICATION USAGE**

6807 These functions are part of the Timers option and need not be available on all implementations.

6808 Note that the absolute value of the monotonic clock is meaningless (because its origin is
 6809 arbitrary), and thus there is no need to set it. Furthermore, realtime applications can rely on the
 6810 fact that the value of this clock is never set and, therefore, that time intervals measured with this
 6811 clock will not be affected by calls to *clock_settime()*.

6812 **RATIONALE**

6813 None.

6814 **FUTURE DIRECTIONS**

6815 None.

6816 **SEE ALSO**

6817 *clock_getcpuclockid()*, *clock_nanosleep()*, *ctime()*, *mq_timedreceive()*, *mq_timedsend()*, *nanosleep()*,
 6818 *pthread_mutex_timedlock()*, *sem_timedwait()*, *time()*, *timer_create()*, *timer_getoverrun()*, the Base
 6819 Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

6820 **CHANGE HISTORY**

6821 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

6822 **Issue 6**

6823 The [ENOSYS] error condition has been removed as stubs need not be provided if an
 6824 implementation does not support the Timers option.

6825 The APPLICATION USAGE section is added.

6826 The following changes were made to align with the IEEE P1003.1a draft standard:

- 6827 • Clarification is added of the effect of resetting the clock resolution.

6828 CPU-time clocks and the *clock_getcpuclockid()* function are added for alignment with
 6829 IEEE Std. 1003.1d-1999.

6830 The following changes are added for alignment with IEEE Std. 1003.1j-2000:

- 6831 • The DESCRIPTION is updated as follows:

- 6832 — The value returned by *clock_gettime()* for CLOCK_MONOTONIC is specified.

- 6833 — *clock_settime()* failing for CLOCK_MONOTONIC is specified.

- 6834 — The effects of *clock_settime()* on the *clock_nanosleep()* function with respect to
 6835 CLOCK_REALTIME is specified.

- 6836 • An [EINVAL] error is added to the ERRORS section, indicating that *clock_settime()* fails for
 6837 CLOCK_MONOTONIC.

- 6838 • The APPLICATION USAGE section notes that the CLOCK_MONOTONIC clock need not
 6839 and shall not be set by *clock_settime()* since the absolute value of the CLOCK_MONOTONIC
 6840 clock is meaningless.

- 6841 • The *clock_nanosleep()*, *mq_timedreceive()*, *mq_timedsend()*, *pthread_mutex_timedlock()*,
 6842 *sem_timedwait()*, *timer_create()*, and *timer_settime()* functions are added to the SEE ALSO
 6843 section.

6844 NAME

6845 clock_nanosleep — high resolution sleep with specifiable clock

6846 SYNOPSIS

6847 cs #include <time.h>

6848 int clock_nanosleep(clockid_t clock_id, int flags,
6849 const struct timespec *rqtp, struct timespec *rmtp);
6850

6851 DESCRIPTION

6852 If the flag `TIMER_ABSTIME` is not set in the *flags* argument, the *clock_nanosleep()* function shall
6853 cause the current thread to be suspended from execution until either the time interval specified
6854 by the *rqtp* argument has elapsed, or a signal is delivered to the calling thread and its action is to
6855 invoke a signal-catching function, or the process is terminated. The clock used to measure the
6856 time shall be the clock specified by *clock_id*.

6857 If the flag `TIMER_ABSTIME` is set in the *flags* argument, the *clock_nanosleep()* function shall
6858 cause the current thread to be suspended from execution until either the time value of the clock
6859 specified by *clock_id* reaches the absolute time specified by the *rqtp* argument, or a signal is
6860 delivered to the calling thread and its action is to invoke a signal-catching function, or the
6861 process is terminated. If, at the time of the call, the time value specified by *rqtp* is less than or
6862 equal to the time value of the specified clock, then *clock_nanosleep()* shall return immediately
6863 and the calling process shall not be suspended.

6864 The suspension time caused by this function may be longer than requested because the
6865 argument value is rounded up to an integer multiple of the sleep resolution, or because of the
6866 scheduling of other activity by the system. But, except for the case of being interrupted by a
6867 signal, the suspension time for the relative *clock_nanosleep()* function (that is, with the
6868 `TIMER_ABSTIME` flag not set) shall not be less than the time interval specified by *rqtp*, as
6869 measured by the corresponding clock. The suspension for the absolute *clock_nanosleep()* function
6870 (that is, with the `TIMER_ABSTIME` flag set) shall be in effect at least until the value of the
6871 corresponding clock reaches the absolute time specified by *rqtp*, except for the case of being
6872 interrupted by a signal.

6873 The use of the *clock_nanosleep()* function shall have no effect on the action or blockage of any
6874 signal.

6875 The *clock_nanosleep()* function shall fail if the *clock_id* argument refers to the CPU-time clock of
6876 the calling thread. It is unspecified if *clock_id* values of other CPU-time clocks are allowed.

6877 RETURN VALUE

6878 If the *clock_nanosleep()* function returns because the requested time has elapsed, its return value
6879 shall be zero.

6880 If the *clock_nanosleep()* function returns because it has been interrupted by a signal, it shall return
6881 the corresponding error value. For the relative *clock_nanosleep()* function, if the *rmtp* argument is
6882 non-NULL, the **timespec** structure referenced by it shall be updated to contain the amount of
6883 time remaining in the interval (the requested time minus the time actually slept). If the *rmtp*
6884 argument is NULL, the remaining time is not returned. The absolute *clock_nanosleep()* function
6885 has no effect on the structure referenced by *rmtp*.

6886 If *clock_nanosleep()* fails, it shall return the corresponding error value.

6887 **ERRORS**6888 The *clock_nanosleep()* function shall fail if:

- 6889 [EINTR] The *clock_nanosleep()* function was interrupted by a signal.
- 6890 [EINVAL] The *rntp* argument specified a nanosecond value less than zero or greater than or equal to 1 000 million; or the `TIMER_ABSTIME` flag was specified in *flags* and the *rntp* argument is outside the range for the clock specified by *clock_id*; or the *clock_id* argument does not specify a known clock, or specifies the CPU-time clock of the calling thread.
- 6895 [ENOTSUP] The *clock_id* argument specifies a clock for which *clock_nanosleep()* is not supported, such as a CPU-time clock.

6897 **EXAMPLES**

6898 None.

6899 **APPLICATION USAGE**

6900 Calling *clock_nanosleep()* with the value `TIMER_ABSTIME` not set in the *flags* argument and with a *clock_id* of `CLOCK_REALTIME` is equivalent to calling *nanosleep()* with the same *rntp* and *rntp* arguments.

6903 **RATIONALE**

6904 The *nanosleep()* function specifies that the system-wide clock `CLOCK_REALTIME` is used to measure the elapsed time for this time service. However, with the introduction of the monotonic clock `CLOCK_MONOTONIC` a new relative sleep function is needed to allow an application to take advantage of the special characteristics of this clock.

6908 There are many applications in which a process needs to be suspended and then activated multiple times in a periodic way; for example, to poll the status of a non-interrupting device or to refresh a display device. For these cases, it is known that precise periodic activation cannot be achieved with a relative *sleep()* or *nanosleep()* function call. Suppose, for example, a periodic process that is activated at time T_0 , executes for a while, and then wants to suspend itself until time T_0+T , the period being T . If this process wants to use the *nanosleep()* function, it must first call *clock_gettime()* to get the current time, then calculate the difference between the current time and T_0+T and, finally, call *nanosleep()* using the computed interval. However, the process could be preempted by a different process between the two function calls, and in this case the interval computed would be wrong; the process would wake up later than desired. This problem would not occur with the absolute *clock_nanosleep()* function, since only one function call would be necessary to suspend the process until the desired time. In other cases, however, a relative sleep is needed, and that is why both functionalities are required.

6921 Although it is possible to implement periodic processes using the timers interface, this implementation would require the use of signals, and the reservation of some signal numbers. In this regard, the reasons for including an absolute version of the *clock_nanosleep()* function in IEEE Std. 1003.1-200x are the same as for the inclusion of the relative *nanosleep()*.

6925 It is also possible to implement precise periodic processes using *pthread_cond_timedwait()*, in which an absolute timeout is specified that takes effect if the condition variable involved is never signaled. However, the use of this interface is unnatural, and involves performing other operations on mutexes and condition variables that imply an unnecessary overhead. Furthermore, *pthread_cond_timedwait()* is not available in implementations that do not support threads.

6931 Although the interface of the relative and absolute versions of the new high resolution sleep service is the same *clock_nanosleep()* function, the *rntp* argument is only used in the relative sleep. This argument is needed in the relative *clock_nanosleep()* function to reissue the function

6934 call if it is interrupted by a signal, but it is not needed in the absolute *clock_nanosleep()* function
6935 call; if the call is interrupted by a signal, the absolute *clock_nanosleep()* function can be invoked
6936 again with the same *rtp* argument used in the interrupted call.

6937 **FUTURE DIRECTIONS**

6938 None.

6939 **SEE ALSO**

6940 *clock_getres()*, *nanosleep()*, *pthread_cond_timedwait()*, *sleep()*, the Base Definitions volume of
6941 IEEE Std. 1003.1-200x, <**time.h**>

6942 **CHANGE HISTORY**

6943 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

6944 **NAME**

6945 clog, clogf, clogl — complex natural logarithm functions

6946 **SYNOPSIS**

6947 #include <complex.h>

6948 double complex clog(double complex *z*);6949 float complex clogf(float complex *z*);6950 long double complex clogl(long double complex *z*);6951 **DESCRIPTION**

6952 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
6953 conflict between the requirements described here and the ISO C standard is unintentional. This
6954 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

6955 These functions shall compute the complex natural (base *e*) logarithm of *z*, with a branch cut
6956 along the negative real axis.

6957 **RETURN VALUE**

6958 These functions shall return the complex natural logarithm value, in the range of a strip
6959 mathematically unbounded along the real axis and in the interval $[-i\pi, +i\pi]$ along the imaginary
6960 axis.

6961 **ERRORS**

6962 No errors are defined.

6963 **EXAMPLES**

6964 None.

6965 **APPLICATION USAGE**

6966 None.

6967 **RATIONALE**

6968 None.

6969 **FUTURE DIRECTIONS**

6970 None.

6971 **SEE ALSO**6972 *cexp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>6973 **CHANGE HISTORY**

6974 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

6975 **NAME**

6976 close — close a file descriptor

6977 **SYNOPSIS**

6978 #include <unistd.h>

6979 int close(int *fildes*);6980 **DESCRIPTION**

6981 The *close()* function shall deallocate the file descriptor indicated by *fildes*. To deallocate means
 6982 to make the file descriptor available for return by subsequent calls to *open()* or other functions
 6983 that allocate file descriptors. All outstanding record locks owned by the process on the file
 6984 associated with the file descriptor shall be removed (that is, unlocked).

6985 If *close()* is interrupted by a signal that is to be caught, it shall return -1 with *errno* set to [EINTR]
 6986 and the state of *fildes* is unspecified. If an I/O error occurred while reading from or writing to the
 6987 file system during *close()*, it may return -1 with *errno* set to [EIO]; if this error is returned, the
 6988 state of *fildes* is unspecified.

6989 When all file descriptors associated with a pipe or FIFO special file are closed, any data
 6990 remaining in the pipe or FIFO shall be discarded.

6991 When all file descriptors associated with an open file description have been closed the open file
 6992 description shall be freed.

6993 If the link count of the file is 0, when all file descriptors associated with the file are closed, the
 6994 space occupied by the file shall be freed and the file shall no longer be accessible.

6995 **XSR** If a STREAMS-based *fildes* is closed and the calling process was previously registered to receive
 6996 a SIGPOLL signal for events associated with that STREAM, the calling process shall be
 6997 unregistered for events associated with the STREAM. The last *close()* for a STREAM causes the
 6998 STREAM associated with *fildes* to be dismantled. If O_NONBLOCK is not set and there have
 6999 been no signals posted for the STREAM, and if there is data on the module's write queue, *close()*
 7000 waits for an unspecified time (for each module and driver) for any output to drain before
 7001 dismantling the STREAM. The time delay can be changed via an I_SETCLTIME *ioctl()* request. If
 7002 the O_NONBLOCK flag is set, or if there are any pending signals, *close()* does not wait for
 7003 output to drain, and dismantles the STREAM immediately.

7004 If the implementation supports STREAMS-based pipes, and *fildes* is associated with one end of a
 7005 pipe, the last *close()* causes a hangup to occur on the other end of the pipe. In addition, if the
 7006 other end of the pipe has been named by *fattach()*, then the last *close()* forces the named end to
 7007 be detached by *fdetach()*. If the named end has no open file descriptors associated with it and
 7008 gets detached, the STREAM associated with that end is also dismantled.

7009 If *fildes* refers to the master side of a pseudo-terminal, and this is the last close, a SIGHUP signal
 7010 is sent to the process group, if any, for which the slave side of the pseudo-terminal is the
 7011 controlling terminal. It is unspecified whether closing the master side of the pseudo-terminal
 7012 flushes all queued input and output.

7013 If *fildes* refers to the slave side of a STREAMS-based pseudo-terminal, a zero-length message
 7014 may be sent to the master.

7015 **AIO** When there is an outstanding cancelable asynchronous I/O operation against *fildes* when *close()*
 7016 is called, that I/O operation may be canceled. An I/O operation that is not canceled completes
 7017 as if the *close()* operation had not yet occurred. All operations that are not canceled shall
 7018 complete as if the *close()* blocked until the operations completed. The *close()* operation itself
 7019 need not block awaiting such I/O completion. Whether any I/O operation is canceled, and
 7020 which I/O operation may be canceled upon *close()*, is implementation-defined.

7021 MF|SHM If a shared memory object or a memory mapped file remains referenced at the last close (that is,
 7022 a process has it mapped), then the entire contents of the memory object shall persist until the
 7023 memory object becomes unreferenced. If this is the last close of a shared memory object or a
 7024 memory mapped file and the close results in the memory object becoming unreferenced, and the
 7025 memory object has been unlinked, then the memory object shall be removed.

7026 If *fdes* refers to a socket, *close()* shall cause the socket to be destroyed. If the socket is in
 7027 connection-mode, and the `SO_LINGER` option is set for the socket with non-zero linger time,
 7028 and the socket has untransmitted data, then *close()* shall block for up to the current linger
 7029 interval until all data is transmitted.

7030 RETURN VALUE

7031 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to
 7032 indicate the error.

7033 ERRORS

7034 The *close()* function shall fail if:

7035 [EBADF] The *fdes* argument is not a valid file descriptor.

7036 [EINTR] The *close()* function was interrupted by a signal.

7037 The *close()* function may fail if:

7038 [EIO] An I/O error occurred while reading from or writing to the file system.

7039 EXAMPLES

7040 Reassigning a File Descriptor

7041 The following example closes the file descriptor associated with standard output for the current
 7042 process, re-assigns standard output to a new file descriptor, and closes the original file
 7043 descriptor to clean up. This example assumes that the file descriptor 0 (which is the descriptor
 7044 for standard input) is not closed.

```
7045 #include <unistd.h>
7046 ...
7047 int pfd;
7048 ...
7049 close(1);
7050 dup(pfd);
7051 close(pfd);
7052 ...
```

7053 Incidentally, this is exactly what could be achieved using:

```
7054 dup2(pfd, 1);
7055 close(pfd);
```

7056 Closing a File Descriptor

7057 In the following example, *close()* is used to close a file descriptor after an unsuccessful attempt is
 7058 made to associate that file descriptor with a stream.

```
7059 #include <stdio.h>
7060 #include <unistd.h>
7061 #include <stdlib.h>
```

```

7062     #define LOCKFILE "/etc/ptmp"
7063     ...
7064     int pfd;
7065     FILE *fpfd;
7066     ...
7067     if ((fpfd = fdopen (pfd, "w")) == NULL) {
7068         close(pfd);
7069         unlink(LOCKFILE);
7070         exit(1);
7071     }
7072     ...

```

7073 APPLICATION USAGE

7074 An application that had used the *stdio* routine *fopen()* to open a file should use the
 7075 corresponding *fclose()* routine rather than *close()*. Once a file is closed, the file descriptor no
 7076 longer exists, since the integer corresponding to it no longer refers to a file.

7077 RATIONALE

7078 The use of interruptible device close routines should be discouraged to avoid problems with the
 7079 implicit closes of file descriptors by *exec* and *exit()*. This volume of IEEE Std. 1003.1-200x only
 7080 intends to permit such behavior by specifying the [EINTR] error condition.

7081 FUTURE DIRECTIONS

7082 None.

7083 SEE ALSO

7084 *fattach()*, *fclose()*, *fdetach()*, *fopen()*, *ioctl()*, *open()*, the Base Definitions volume of
 7085 IEEE Std. 1003.1-200x, <**unistd.h**>, Section 2.6 (on page 539)

7086 CHANGE HISTORY

7087 First released in Issue 1. Derived from Issue 1 of the SVID.

7088 Issue 4

7089 The <**unistd.h**> header is added to the SYNOPSIS section.

7090 Issue 4, Version 2

7091 The following changes are incorporated for X/OPEN UNIX conformance:

- 7092 • The DESCRIPTION is updated to describe the actions of closing a file descriptor referring to
 7093 a STREAMS-based file or either side of a pseudo-terminal.
- 7094 • The ERRORS section describes a condition under which the [EIO] error may be returned.

7095 Issue 5

7096 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.

7097 Issue 6

7098 The DESCRIPTION related to a STREAMS-based file or pseudo-terminal is marked as part of the
 7099 XSI STREAMS Option Group.

7100 The following new requirements on POSIX implementations derive from alignment with the
 7101 Single UNIX Specification:

- 7102 • The [EIO] error condition is added as an optional error.
- 7103 • The DESCRIPTION is updated to describe the state of the *fildev* file descriptor as unspecified
 7104 if an I/O error occurs and an [EIO] error condition is returned.

7105 Text referring to sockets is added to the DESCRIPTION.

7106
7107
7108

The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that shared memory objects and memory mapped files (and not typed memory objects) are the types of memory objects to which the paragraph on last closes applies.

7109 **NAME**

7110 closedir — close a directory stream

7111 **SYNOPSIS**

7112 #include <dirent.h>

7113 int closedir(DIR *dirp);

7114 **DESCRIPTION**

7115 The *closedir()* function shall close the directory stream referred to by the argument *dirp*. Upon
7116 return, the value of *dirp* may no longer point to an accessible object of the type **DIR**. If a file
7117 descriptor is used to implement type **DIR**, that file descriptor shall be closed.

7118 **RETURN VALUE**

7119 Upon successful completion, *closedir()* shall return 0; otherwise, -1 shall be returned and *errno*
7120 set to indicate the error.

7121 **ERRORS**7122 The *closedir()* function may fail if:7123 [EBADF] The *dirp* argument does not refer to an open directory stream. |7124 [EINTR] The *closedir()* function was interrupted by a signal. |7125 **EXAMPLES**7126 **Closing a Directory Stream**7127 The following program fragment demonstrates how the *closedir()* function is used.

```
7128 ...  
7129     DIR *dir;  
7130     struct dirent *dp;  
7131 ...  
7132     if ((dir = opendir(".")) == NULL) {  
7133 ...  
7134     }  
7135     while ((dp = readdir(dir)) != NULL) {  
7136 ...  
7137     }  
7138     closedir(dir);  
7139 ...
```

7140 **APPLICATION USAGE**

7141 None.

7142 **RATIONALE**

7143 None.

7144 **FUTURE DIRECTIONS**

7145 None.

7146 **SEE ALSO**7147 *opendir()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <dirent.h> |

7148 **CHANGE HISTORY**

7149 First released in Issue 2.

7150 **Issue 4**7151 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on
7152 XSI-conformant systems.

7153 The [EINTR] error is marked as an extension.

7154 **Issue 6**7155 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.7156 The following new requirements on POSIX implementations derive from alignment with the
7157 Single UNIX Specification:

- 7158 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was
7159 required for conforming implementations of previous POSIX specifications, it was not
7160 required for UNIX applications.
- 7161 • The [EINTR] error condition is added as an optional error condition.

7162 NAME

7163 closelog, openlog, setlogmask, syslog — control system log

7164 SYNOPSIS

```

7165 xSI #include <syslog.h>
7166
7166 void closelog(void);
7167 void openlog(const char *ident, int logopt, int facility);
7168 int setlogmask(int maskpri);
7169 void syslog(int priority, const char *message, ... /* arguments */);
7170

```

7171 DESCRIPTION

7172 The *syslog()* function shall send a message to an implementation-defined logging facility, which
 7173 may log it in an implementation-defined system log, write it to the system console, forward it to
 7174 a list of users, or forward it to the logging facility on another host over the network. The logged
 7175 message shall include a message header and a message body. The message header contains at
 7176 least a timestamp and a tag string.

7177 The message body is generated from the *message* and following arguments in the same manner
 7178 as if these were arguments to *printf()*, except that occurrences of %m in the format string
 7179 pointed to by the *message* argument are replaced by the error message string associated with the
 7180 current value of *errno*. A trailing <newline> character is added if needed.

7181 Values of the *priority* argument are formed by OR'ing together a severity level value and an
 7182 optional facility value. If no facility value is specified, the current default facility value is used.

7183 Possible values of severity level include:

7184	LOG_EMERG	A panic condition.
7185	LOG_ALERT	A condition that should be corrected immediately, such as a corrupted system database.
7186		
7187	LOG_CRIT	Critical conditions, such as hard device errors.
7188	LOG_ERR	Errors.
7189	LOG_WARNING	
7190		Warning messages.
7191	LOG_NOTICE	Conditions that are not error conditions, but that may require special handling.
7192		
7193	LOG_INFO	Informational messages.
7194	LOG_DEBUG	Messages that contain information normally of use only when debugging a program.
7195		

7196 The facility indicates the application or system component generating the message. Possible
 7197 facility values include:

7198	LOG_USER	Messages generated by arbitrary processes. This is the default facility identifier if none is specified.
7199		
7200	LOG_LOCAL0	Reserved for local use.
7201	LOG_LOCAL1	Reserved for local use.
7202	LOG_LOCAL2	Reserved for local use.

- 7203 LOG_LOCAL3 Reserved for local use.
- 7204 LOG_LOCAL4 Reserved for local use.
- 7205 LOG_LOCAL5 Reserved for local use.
- 7206 LOG_LOCAL6 Reserved for local use.
- 7207 LOG_LOCAL7 Reserved for local use.
- 7208 The *openlog()* function shall set process attributes that affect subsequent calls to *syslog()*. The
- 7209 *ident* argument is a string that is prepended to every message. The *logopt* argument indicates
- 7210 logging options. Values for *logopt* are constructed by a bitwise-inclusive OR of zero or more of
- 7211 the following:
- 7212 LOG_PID Log the process ID with each message. This is useful for identifying specific
- 7213 processes.
- 7214 LOG_CONS Write messages to the system console if they cannot be sent to the logging
- 7215 facility. The *syslog()* function ensures that the process does not acquire the
- 7216 console as a controlling terminal in the process of writing the message.
- 7217 LOG_NDELAY Open the connection to the logging facility immediately. Normally the open is
- 7218 delayed until the first message is logged. This is useful for programs that need
- 7219 to manage the order in which file descriptors are allocated.
- 7220 LOG_ODELAY Delay open until *syslog()* is called.
- 7221 LOG_NOWAIT Do not wait for child processes that may have been created during the course
- 7222 of logging the message. This option should be used by processes that enable
- 7223 notification of child termination using SIGCHLD, since *syslog()* may
- 7224 otherwise block waiting for a child whose exit status has already been
- 7225 collected.
- 7226 The *facility* argument encodes a default facility to be assigned to all messages that do not have
- 7227 an explicit facility already encoded. The initial default facility is LOG_USER.
- 7228 The *openlog()* and *syslog()* functions may allocate a file descriptor. It is not necessary to call
- 7229 *openlog()* prior to calling *syslog()*.
- 7230 The *closelog()* function shall close any open file descriptors allocated by previous calls to
- 7231 *openlog()* or *syslog()*.
- 7232 The *setlogmask()* function shall set the log priority mask for the current process to *maskpri* and
- 7233 return the previous mask. If the *maskpri* argument is 0, the current log mask is not modified.
- 7234 Calls by the current process to *syslog()* with a priority not set in *maskpri* shall be rejected. The
- 7235 default log mask allows all priorities to be logged. A call to *openlog()* is not required prior to
- 7236 calling *setlogmask()*.
- 7237 Symbolic constants for use as values of the *logopt*, *facility*, *priority*, and *maskpri* arguments are
- 7238 defined in the `<syslog.h>` header.
- 7239 **RETURN VALUE**
- 7240 The *setlogmask()* function shall return the previous log priority mask. The *closelog()*, *openlog()*,
- 7241 and *syslog()* functions shall return no value.
- 7242 **ERRORS**
- 7243 No errors are defined.

7244 **EXAMPLES**7245 **Using openlog()**

7246 The following example causes subsequent calls to *syslog()* to log the process ID with each
7247 message, and to write messages to the system console if they cannot be sent to the logging
7248 facility.

```
7249 #include <syslog.h>
7250 char *ident = "Process demo";
7251 int logopt = LOG_PID | LOG_CONS;
7252 int facility = LOG_USER;
7253 ...
7254 openlog(ident, logopt, facility);
```

7255 **Using setlogmask()**

7256 The following example causes subsequent calls to *syslog()* to accept error messages or messages
7257 generated by arbitrary processes, and to reject all other messages.

```
7258 #include <syslog.h>
7259 int result;
7260 int mask = LOG_MASK (LOG_ERR | LOG_USER);
7261 ...
7262 result = setlogmask(mask);
```

7263 **Using syslog**

7264 The following example sends the message "This is a message" to the default logging
7265 facility, marking the message as an error message generated by random processes.

```
7266 #include <syslog.h>
7267 char *message = "This is a message";
7268 int priority = LOG_ERR | LOG_USER;
7269 ...
7270 syslog(priority, message);
```

7271 **APPLICATION USAGE**

7272 None.

7273 **RATIONALE**

7274 None.

7275 **FUTURE DIRECTIONS**

7276 None.

7277 **SEE ALSO**

7278 *printf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**syslog.h**>

7279 **CHANGE HISTORY**

7280 First released in Issue 4, Version 2.

7281 **Issue 5**

7282 Moved from X/OPEN UNIX extension to BASE.

7283 NAME

7284 confstr — get configurable variables

7285 SYNOPSIS

7286 #include <unistd.h>

7287 size_t confstr(int name, char *buf, size_t len);

7288 DESCRIPTION

7289 The *confstr()* function provides a method for applications to get configuration-defined string
 7290 values. Its use and purpose are similar to *sysconf()*, but it is used where string values rather than
 7291 numeric values are returned.

7292 The *name* argument represents the system variable to be queried. The implementation shall
 7293 support the following name values, defined in <unistd.h>. It may support others:

7294 _CS_PATH
 7295 _CS_POSIX_V6_ILP32_OFF32_CFLAGS
 7296 _CS_POSIX_V6_ILP32_OFF32_LDFLAGS
 7297 _CS_POSIX_V6_ILP32_OFF32_LIBS
 7298 _CS_POSIX_V6_ILP32_OFF32_LINTFLAGS
 7299 _CS_POSIX_V6_ILP32_OFFBIG_CFLAGS
 7300 _CS_POSIX_V6_ILP32_OFFBIG_LDFLAGS
 7301 _CS_POSIX_V6_ILP32_OFFBIG_LIBS
 7302 _CS_POSIX_V6_ILP32_OFFBIG_LINTFLAGS
 7303 _CS_POSIX_V6_LP64_OFF64_CFLAGS
 7304 _CS_POSIX_V6_LP64_OFF64_LDFLAGS
 7305 _CS_POSIX_V6_LP64_OFF64_LIBS
 7306 _CS_POSIX_V6_LP64_OFF64_LINTFLAGS
 7307 _CS_POSIX_V6_LPBIG_OFFBIG_CFLAGS
 7308 _CS_POSIX_V6_LPBIG_OFFBIG_LDFLAGS
 7309 _CS_POSIX_V6_LPBIG_OFFBIG_LIBS
 7310 _CS_POSIX_V6_LPBIG_OFFBIG_LINTFLAGS
 7311 XSI CS_XBS5_ILP32_OFF32_CFLAGS (LEGACY)
 7312 CS_XBS5_ILP32_OFF32_LDFLAGS (LEGACY)
 7313 CS_XBS5_ILP32_OFF32_LIBS (LEGACY)
 7314 CS_XBS5_ILP32_OFF32_LINTFLAGS (LEGACY)
 7315 CS_XBS5_ILP32_OFFBIG_CFLAGS (LEGACY)
 7316 CS_XBS5_ILP32_OFFBIG_LDFLAGS (LEGACY)
 7317 CS_XBS5_ILP32_OFFBIG_LIBS (LEGACY)
 7318 CS_XBS5_ILP32_OFFBIG_LINTFLAGS (LEGACY)
 7319 CS_XBS5_LP64_OFF64_CFLAGS (LEGACY)
 7320 CS_XBS5_LP64_OFF64_LDFLAGS (LEGACY)
 7321 CS_XBS5_LP64_OFF64_LIBS (LEGACY)
 7322 CS_XBS5_LP64_OFF64_LINTFLAGS (LEGACY)
 7323 CS_XBS5_LPBIG_OFFBIG_CFLAGS (LEGACY)
 7324 CS_XBS5_LPBIG_OFFBIG_LDFLAGS (LEGACY)
 7325 CS_XBS5_LPBIG_OFFBIG_LIBS (LEGACY)
 7326 CS_XBS5_LPBIG_OFFBIG_LINTFLAGS (LEGACY)

7327

7328 If *len* is not 0, and if *name* has a configuration-defined value, *confstr()* shall copy that value into
 7329 the *len*-byte buffer pointed to by *buf*. If the string to be returned is longer than *len* bytes,
 7330 including the terminating null, then *confstr()* shall truncate the string to *len*-1 bytes and null-
 7331 terminate the result. The application can detect that the string was truncated by comparing the

7332 value returned by *confstr()* with *len*.

7333 If *len* is 0 and *buf* is a null pointer, then *confstr()* shall still return the integer value as defined
7334 below, but shall not return a string. If *len* is 0 but *buf* is not a null pointer, the result is
7335 unspecified.

7336 If the implementation supports the Shell option, the string stored in *buf* after a call to:

```
7337 confstr(_CS_PATH, buf, sizeof(buf))
```

7338 can be used as a value of the *PATH* environment variable that accesses all of the standard
7339 utilities of IEEE Std. 1003.1-200x, if the return value is less than or equal to *sizeof(buf)*.

7340 RETURN VALUE

7341 If *name* has a configuration-defined value, *confstr()* shall return the size of buffer that would be
7342 needed to hold the entire configuration-defined value including the terminating null. If this
7343 return value is greater than *len*, the string returned in *buf* is truncated.

7344 If *name* is invalid, *confstr()* shall return 0 and set *errno* to indicate the error.

7345 If *name* does not have a configuration-defined value, *confstr()* shall return 0 and leave *errno*
7346 unchanged.

7347 ERRORS

7348 The *confstr()* function shall fail if:

7349 [EINVAL] The value of the *name* argument is invalid.

7350 EXAMPLES

7351 None.

7352 APPLICATION USAGE

7353 An application can distinguish between an invalid *name* parameter value and one that
7354 corresponds to a configurable variable that has no configuration-defined value by checking if
7355 *errno* is modified. This mirrors the behavior of *sysconf()*.

7356 The original need for this function was to provide a way of finding the configuration-defined
7357 default value for the environment variable *PATH*. Since *PATH* can be modified by the user to
7358 include directories that could contain utilities replacing the standard utilities in the Shell and
7359 Utilities volume of IEEE Std. 1003.1-200x, applications need a way to determine the system-
7360 supplied *PATH* environment variable value that contains the correct search path for the standard
7361 utilities.

7362 An application could use:

```
7363 confstr(name, (char *)NULL, (size_t)0)
```

7364 to find out how big a buffer is needed for the string value; use *malloc()* to allocate a buffer to
7365 hold the string; and call *confstr()* again to get the string. Alternately, it could allocate a fixed,
7366 static buffer that is big enough to hold most answers (perhaps 512 or 1024 bytes), but then use
7367 *malloc()* to allocate a larger buffer if it finds that this is too small.

7368 RATIONALE

7369 Application developers can normally determine any configuration variable by means of reading
7370 from the stream opened by a call to:

```
7371 popen("command -p getconf variable", "r");
```

7372 The *confstr()* function with a *name* argument of *_CS_PATH* returns a string that can be used as a
7373 *PATH* environment variable setting that will reference the standard shell and utilities as
7374 described in the Shell and Utilities volume of IEEE Std. 1003.1-200x.

7375 The *confstr()* function copies the returned string into a buffer supplied by the application instead
7376 of returning a pointer to a string. This allows a cleaner function in some implementations (such
7377 as those with lightweight threads) and resolves questions about when the application must copy
7378 the string returned.

7379 **FUTURE DIRECTIONS**

7380 None.

7381 **SEE ALSO**

7382 *pathconf()*, *sysconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>, the Shell
7383 and Utilities volume of IEEE Std. 1003.1-200x, *c99*

7384 **CHANGE HISTORY**

7385 First released in Issue 4. Derived from the ISO POSIX-2 standard.

7386 **Issue 5**

7387 A table indicating the permissible values of *name* are added to the DESCRIPTION. All those
7388 marked EX are new in this issue.

7389 **Issue 6**

7390 The Open Group corrigenda item U033/7 has been applied. The return value for the case
7391 returning the size of the buffer now explicitly states that this includes the terminating null.

7392 The following new requirements on POSIX implementations derive from alignment with the
7393 Single UNIX Specification:

- 7394 • The DESCRIPTION is updated with new arguments which can be used to determine
7395 configuration strings for C compiler flags, linker/loader flags, and libraries for each different
7396 supported programming environment. This is a change to support data size neutrality.

7397 The following changes were made to align with the IEEE P1003.1a draft standard:

- 7398 • The DESCRIPTION is updated to include text describing how `_CS_PATH` can be used to
7399 obtain a *PATH* to access the standard utilities.

7400 The macros associated with the *c89* programming models are marked LEGACY and new
7401 equivalent macros associated with *c99* are introduced.

7402 **NAME**

7403 conj, conjf, conjl — complex conjugate functions

7404 **SYNOPSIS**

7405 #include <complex.h>

7406 double complex conj(double complex *z*);7407 float complex conjf(float complex *z*);7408 long double complex conjl(long double complex *z*);7409 **DESCRIPTION**

7410 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
7411 conflict between the requirements described here and the ISO C standard is unintentional. This
7412 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

7413 These functions shall compute the complex conjugate of *z*, by reversing the sign of its imaginary
7414 part.

7415 **RETURN VALUE**

7416 These functions return the complex conjugate value.

7417 **ERRORS**

7418 No errors are defined.

7419 **EXAMPLES**

7420 None.

7421 **APPLICATION USAGE**

7422 None.

7423 **RATIONALE**

7424 None.

7425 **FUTURE DIRECTIONS**

7426 None.

7427 **SEE ALSO**

7428 *carg()*, *cimag()*, *cproj()*, *creal()*, the Base Definitions volume of IEEE Std. 1003.1-200x,
7429 <complex.h>

7430 **CHANGE HISTORY**

7431 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7432 **NAME**

7433 connect — connect a socket

7434 **SYNOPSIS**

7435 #include <sys/socket.h>

7436 int connect(int *socket*, const struct sockaddr **address*,
7437 socklen_t *address_len*);7438 **DESCRIPTION**7439 The *connect()* function requests a connection to be made on a socket. The function takes the
7440 following arguments:

7441	<i>socket</i>	Specifies the file descriptor associated with the socket.
7442	<i>address</i>	Points to a sockaddr structure containing the peer address. The length and 7443 format of the address depend on the address family of the socket.
7444	<i>address_len</i>	Specifies the length of the sockaddr structure pointed to by the <i>address</i> 7445 argument.

7446 If the socket has not already been bound to a local address, *connect()* shall bind it to an address
7447 which, unless the socket's address family is AF_UNIX, is an unused local address.7448 If the initiating socket is not connection-mode, then *connect()* shall set the socket's peer address,
7449 and no connection is made. For SOCK_DGRAM sockets, the peer address identifies where all
7450 datagrams are sent on subsequent *send()* functions, and limits the remote sender for subsequent
7451 *recv()* functions. If *address* is a null address for the protocol, the socket's peer address shall be
7452 reset.7453 If the initiating socket is connection-mode, then *connect()* attempts to establish a connection to
7454 the address specified by the *address* argument.7455 If the connection cannot be established immediately and O_NONBLOCK is not set for the file
7456 descriptor for the socket, *connect()* shall block for up to an unspecified timeout interval until the
7457 connection is established. If the timeout interval expires before the connection is established,
7458 *connect()* shall fail and the connection attempt shall be aborted. If *connect()* is interrupted by a
7459 signal that is caught while blocked waiting to establish a connection, *connect()* shall fail and set
7460 *errno* to [EINTR], but the connection request shall not be aborted, and the connection shall be
7461 established asynchronously.7462 If the connection cannot be established immediately and O_NONBLOCK is set for the file
7463 descriptor for the socket, *connect()* shall fail and set *errno* to [EINPROGRESS], but the connection
7464 request shall not be aborted, and the connection shall be established asynchronously.
7465 Subsequent calls to *connect()* for the same socket, before the connection is established, shall fail
7466 and set *errno* to [EALREADY].7467 When the connection has been established asynchronously, *select()* and *poll()* shall indicate that
7468 the file descriptor for the socket is ready for writing.7469 The socket in use may require the process to have appropriate privileges to use the *connect()*
7470 function.7471 **RETURN VALUE**7472 Upon successful completion, *connect()* shall return 0; otherwise, -1 shall be returned and *errno*
7473 set to indicate the error.

7474 **ERRORS**

- 7475 The *connect()* function shall fail if:
- 7476 [EADDRNOTAVAIL]
7477 The specified address is not available from the local machine.
- 7478 [EAFNOSUPPORT]
7479 The specified address is not a valid address for the address family of the
7480 specified socket.
- 7481 [EALREADY] A connection request is already in progress for the specified socket.
- 7482 [EBADF] The *socket* argument is not a valid file descriptor.
- 7483 [ECONNREFUSED]
7484 The target address was not listening for connections or refused the connection
7485 request.
- 7486 [EINPROGRESS] O_NONBLOCK is set for the file descriptor for the socket and the connection
7487 cannot be immediately established; the connection shall be established
7488 asynchronously.
- 7489 [EINTR] The attempt to establish a connection was interrupted by delivery of a signal
7490 that was caught; the connection shall be established asynchronously.
- 7491 [EISCONN] The specified socket is connection-mode and is already connected.
- 7492 [ENETUNREACH]
7493 No route to the network is present.
- 7494 [ENOTSOCK] The *socket* argument does not refer to a socket.
- 7495 [EPROTOTYPE] The specified address has a different type than the socket bound to the
7496 specified peer address.
- 7497 [ETIMEDOUT] The attempt to connect timed out before a connection was made.
- 7498 If the address family of the socket is AF_UNIX, then *connect()* shall fail if:
- 7499 [EIO] An I/O error occurred while reading from or writing to the file system.
- 7500 [ELOOP] A loop exists in symbolic links encountered during resolution of the path
7501 name in *address*.
- 7502 [ENAMETOOLONG]
7503 A component of a path name exceeded {NAME_MAX} characters, or an entire
7504 path name exceeded {PATH_MAX} characters.
- 7505 [ENOENT] A component of the path name does not name an existing file or the path
7506 name is an empty string.
- 7507 [ENOTDIR] A component of the path prefix of the path name in *address* is not a directory.
- 7508 The *connect()* function may fail if:
- 7509 [EACCES] Search permission is denied for a component of the path prefix; or write
7510 access to the named socket is denied.
- 7511 [EADDRINUSE] Attempt to establish a connection that uses addresses that are already in use.
- 7512 [ECONNRESET] Remote host reset the connection request.
- 7513 [EHOSTUNREACH]
7514 The destination host cannot be reached (probably because the host is down or

7515 a remote router cannot reach it).

7516 [EINVAL] The *address_len* argument is not a valid length for the address family; or
7517 invalid address family in the **sockaddr** structure.

7518 [ELOOP] More than {SYMLOOP_MAX} symbolic links were encountered during
7519 resolution of the path name in *address*.

7520 [ENAMETOOLONG]
7521 Path name resolution of a symbolic link produced an intermediate result
7522 whose length exceeds {PATH_MAX}.

7523 [ENETDOWN] The local network interface used to reach the destination is down.

7524 [ENOBUFS] No buffer space is available.

7525 [EOPNOTSUPP] The socket is listening and cannot be connected.

7526 **EXAMPLES**

7527 None.

7528 **APPLICATION USAGE**

7529 If *connect()* fails, the state of the socket is unspecified. Portable applications should close the file
7530 descriptor and create a new socket before attempting to reconnect.

7531 **RATIONALE**

7532 None.

7533 **FUTURE DIRECTIONS**

7534 None.

7535 **SEE ALSO**

7536 *accept()*, *bind()*, *close()*, *getsockname()*, *poll()*, *select()*, *send()*, *shutdown()*, *socket()*, the Base
7537 Definitions volume of IEEE Std. 1003.1-200x, <sys/socket.h>

7538 **CHANGE HISTORY**

7539 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

7540 The wording of the mandatory [ELOOP] error condition is updated, and a second optional
7541 [ELOOP] error condition is added.

7542 **NAME**

7543 copysign, copysignf, copysignl — number manipulation function

7544 **SYNOPSIS**

7545 #include <math.h>

7546 double copysign(double x, double y);

7547 float copysignf(float x, float y);

7548 long double copysignl(long double x, long double y);

7549 **DESCRIPTION**

7550 cx The functionality described on this reference page is aligned with the ISO C standard. Any
7551 conflict between the requirements described here and the ISO C standard is unintentional. This
7552 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

7553 These functions shall produce a value with the magnitude of *x* and the sign of *y*. They produce a
7554 NaN (with the sign of *y*) if *x* is a NaN. On implementations that represent a signed zero but do
7555 not treat negative zero consistently in arithmetic operations, these functions regard the sign of
7556 zero as positive.

7557 An application wishing to check for error situations should set *errno* to 0 before calling these
7558 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

7559 **RETURN VALUE**

7560 Upon successful completion, these functions shall return a value with the magnitude of *x* and
7561 the sign of *y*.

7562 If *x* is $\pm\text{Inf}$, these functions shall return *x*.

7563 If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

7564 **ERRORS**

7565 These functions may fail if:

7566 [EDOM] The value of *x* is NaN.

7567 **EXAMPLES**

7568 None.

7569 **APPLICATION USAGE**

7570 None.

7571 **RATIONALE**

7572 *copysign()* and *signbit()* need not be consistent with each other if the arithmetic is not consistent
7573 in its treatment of zeros. For example, the IBM S/370 has instructions to flip the sign bit making
7574 it possible to create a negative zero, but $\pm 0.0 \times \pm 1.0$ is always $+0.0$. In this case, *copysign()* will
7575 treat 0.0 as positive, while *signbit()* will treat it as negative.

7576 **FUTURE DIRECTIONS**

7577 None.

7578 **SEE ALSO**

7579 *signbit()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

7580 **CHANGE HISTORY**

7581 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7582 **NAME**

7583 cos, cosf, cosl — cosine function

7584 **SYNOPSIS**

7585 #include <math.h>

7586 double cos(double x);

7587 float cosf(float x);

7588 long double cosl(long double x);

7589 **DESCRIPTION**

7590 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
 7591 conflict between the requirements described here and the ISO C standard is unintentional. This
 7592 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

7593 These functions shall compute the cosine of x , measured in radians.

7594 An application wishing to check for error situations should set *errno* to 0 before calling *cos()*. If
 7595 *errno* is non-zero on return, or the returned value is NaN, an error has occurred.

7596 **RETURN VALUE**7597 Upon successful completion, these functions shall return the cosine of x .7598 **XSI** If x is NaN, NaN shall be returned and *errno* may be set to [EDOM].

7599 **XSI** If x is $\pm\text{Inf}$, either 0 shall be returned and *errno* set to [EDOM], or NaN shall be returned and *errno*
 7600 may be set to [EDOM].

7601 If the result underflows, 0 shall be returned and *errno* may be set to [ERANGE].7602 **ERRORS**

7603 These functions may fail if:

7604 **XSI** [EDOM] The value of x is NaN or x is $\pm\text{Inf}$.

7605 [ERANGE] The result underflows

7606 **XSI** No other errors shall occur.7607 **EXAMPLES**7608 **Taking the Cosine of a 45-Degree Angle**

7609 #include <math.h>

7610 ...

7611 double radians = 45 * M_PI / 180;

7612 double result;

7613 ...

7614 result = cos(radians);

7615 **APPLICATION USAGE**7616 The *cos()* function may lose accuracy when its argument is far from 0.7617 **RATIONALE**

7618 None.

7619 **FUTURE DIRECTIONS**

7620 None.

7621 **SEE ALSO**

7622 *acos()*, *isnan()*, *sin()*, *tan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

7623 **CHANGE HISTORY**

7624 First released in Issue 1. Derived from Issue 1 of the SVID.

7625 **Issue 4**

7626 References to *matherr()* are removed.

7627 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the ISO C standard and to rationalize error handling in the mathematics functions.

7629 The return value specified for [EDOM] is marked as an extension.

7630 **Issue 5**

7631 The DESCRIPTION is updated to indicate how an application should check for an error. This text was previously published in the APPLICATION USAGE section.

7633 **Issue 6**

7634 The *cosf()* and *cosl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

7635 **NAME**

7636 cosh, coshf, coshl — hyperbolic cosine function

7637 **SYNOPSIS**

7638 #include <math.h>

7639 double cosh(double x);

7640 float coshf(float x);

7641 long double coshl(long double x);

7642 **DESCRIPTION**7643 cx The functionality described on this reference page is aligned with the ISO C standard. Any
7644 conflict between the requirements described here and the ISO C standard is unintentional. This
7645 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.7646 These functions shall compute the hyperbolic cosine of x .7647 An application wishing to check for error situations should set *errno* to 0 before calling *cosh()*. If
7648 *errno* is non-zero on return, or the returned value is NaN, an error has occurred.7649 **RETURN VALUE**7650 Upon successful completion, these functions shall return the hyperbolic cosine of x .7651 If the result would cause an overflow, HUGE_VAL shall be returned and *errno* set to [ERANGE].7652 xSI If x is NaN, NaN shall be returned and *errno* may be set to [EDOM].7653 **ERRORS**

7654 These functions shall fail if:

7655 [ERANGE] The result would cause an overflow.

7656 These functions may fail if:

7657 xSI [EDOM] The value of x is NaN.

7658 xSI No other errors shall occur.

7659 **EXAMPLES**

7660 None.

7661 **APPLICATION USAGE**

7662 None.

7663 **RATIONALE**

7664 None.

7665 **FUTURE DIRECTIONS**

7666 None.

7667 **SEE ALSO**7668 *acosh()*, *isnan()*, *sinh()*, *tanh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>7669 **CHANGE HISTORY**

7670 First released in Issue 1. Derived from Issue 1 of the SVID.

7671 **Issue 4**7672 References to *matherr()* are removed.7673 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the
7674 ISO C standard and to rationalize error handling in the mathematics functions.

- 7675 The return value specified for [EDOM] is marked as an extension.
- 7676 **Issue 5**
- 7677 The DESCRIPTION is updated to indicate how an application should check for an error. This
- 7678 text was previously published in the APPLICATION USAGE section.
- 7679 **Issue 6**
- 7680 The *coshf()* and *coshl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

7681 **NAME**

7682 cpow, cpowf, cpowl — complex power functions

7683 **SYNOPSIS**

7684 #include <complex.h>

7685 double complex cpow(double complex *x*, double complex *y*);7686 float complex cpowf(float complex *x*, float complex *y*);7687 long double complex cpowl(long double complex *x*,7688 long double complex *y*);7689 **DESCRIPTION**7690 *CX* The functionality described on this reference page is aligned with the ISO C standard. Any
7691 conflict between the requirements described here and the ISO C standard is unintentional. This
7692 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.7693 These functions shall compute the complex power function x^y , with a branch cut for the first
7694 parameter along the negative real axis.7695 **RETURN VALUE**

7696 These functions shall return the complex power function value.

7697 **ERRORS**

7698 No errors are defined.

7699 **EXAMPLES**

7700 None.

7701 **APPLICATION USAGE**

7702 None.

7703 **RATIONALE**

7704 None.

7705 **FUTURE DIRECTIONS**

7706 None.

7707 **SEE ALSO**7708 *cabs()*, *csqrt()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>7709 **CHANGE HISTORY**

7710 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7711 **NAME**

7712 cproj, cprojf, cprojl — complex projection functions

7713 **SYNOPSIS**

7714 #include <complex.h>

7715 double complex cproj(double complex z);

7716 float complex cprojf(float complex z);

7717 long double complex cprojl(long double complex z);

7718 **DESCRIPTION**

7719 cx The functionality described on this reference page is aligned with the ISO C standard. Any
7720 conflict between the requirements described here and the ISO C standard is unintentional. This
7721 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

7722 These functions shall compute a projection of z onto the Riemann sphere: z projects to z , except
7723 that all complex infinities (even those with one infinite part and one NaN part) project to
7724 positive infinity on the real axis. If z has an infinite part, then $cproj(z)$ is equivalent to:

7725 $\text{INFINITY} + \text{I} * \text{copysign}(0.0, \text{cimag}(z))$ 7726 **RETURN VALUE**

7727 These functions shall return the value of the projection onto the Riemann sphere.

7728 **ERRORS**

7729 No errors are defined.

7730 **EXAMPLES**

7731 None.

7732 **APPLICATION USAGE**

7733 None.

7734 **RATIONALE**

7735 Two topologies are commonly used in complex mathematics: the complex plane with its
7736 continuum of infinities, and the Riemann sphere with its single infinity. The complex plane is
7737 better suited for transcendental functions, the Riemann sphere for algebraic functions. The
7738 complex types with their multiplicity of infinities provide a useful (though imperfect) model for
7739 the complex plane. The $cproj()$ function helps model the Riemann sphere by mapping all
7740 infinities to one, and should be used just before any operation, especially comparisons, that
7741 might give spurious results for any of the other infinities. Note that a complex value with one
7742 infinite part and one NaN part is regarded as an infinity, not a NaN, because if one part is
7743 infinite, the complex value is infinite independent of the value of the other part. For the same
7744 reason, $cabs()$ returns an infinity if its argument has an infinite part and a NaN part.

7745 **FUTURE DIRECTIONS**

7746 None.

7747 **SEE ALSO**

7748 $carg()$, $cimag()$, $conj()$, $creal()$, the Base Definitions volume of IEEE Std. 1003.1-200x,
7749 <complex.h>

7750 **CHANGE HISTORY**

7751 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7752 **NAME**

7753 creal, crealf, creall — complex real functions

7754 **SYNOPSIS**

7755 #include <complex.h>

7756 double creal(double complex z);

7757 float crealf(float complex z);

7758 long double creall(long double complex z);

7759 **DESCRIPTION**

7760 cx The functionality described on this reference page is aligned with the ISO C standard. Any
7761 conflict between the requirements described here and the ISO C standard is unintentional. This
7762 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

7763 These functions shall compute the real part of *z*.7764 **RETURN VALUE**

7765 These functions shall return the real part value.

7766 **ERRORS**

7767 No errors are defined.

7768 **EXAMPLES**

7769 None.

7770 **APPLICATION USAGE**7771 For a variable *z* of complex type:7772 `z == creal(z) + cimag(z)*I`7773 **RATIONALE**

7774 None.

7775 **FUTURE DIRECTIONS**

7776 None.

7777 **SEE ALSO**

7778 *carg()*, *cimag()*, *conj()*, *cproj()*, the Base Definitions volume of IEEE Std. 1003.1-200x,
7779 <complex.h>

7780 **CHANGE HISTORY**

7781 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7782 **NAME**

7783 creat — create a new file or rewrite an existing one

7784 **SYNOPSIS**

7785 OH #include <sys/stat.h>

7786 #include <fcntl.h>

7787 int creat(const char *path, mode_t mode);

7788 **DESCRIPTION**

7789 The function call:

7790 creat(path, mode)

7791 is equivalent to:

7792 open(path, O_WRONLY|O_CREAT|O_TRUNC, mode)

7793 **RETURN VALUE**7794 Refer to *open()*.7795 **ERRORS**7796 Refer to *open()*.7797 **EXAMPLES**7798 **Creating a File**7799 The following example creates the file **/tmp/file** with read and write permissions for the file
7800 owner and read permission for group and others. The resulting file descriptor is assigned to the
7801 *fd* variable.

7802 #include <fcntl.h>

7803 ...

7804 int fd;

7805 mode_t mode = S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH;

7806 char *filename = "/tmp/file";

7807 ...

7808 fd = creat(filename, mode);

7809 ...

7810 **APPLICATION USAGE**

7811 None.

7812 **RATIONALE**7813 The *creat()* function is redundant. Its services are also provided by the *open()* function. It has
7814 been included primarily for historical purposes since many existing applications depend on it. It
7815 is best considered a part of the C binding rather than a function that should be provided in other
7816 languages.7817 **FUTURE DIRECTIONS**

7818 None.

7819 **SEE ALSO**7820 *open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <fcntl.h>, <sys/stat.h>, |
7821 <sys/types.h>

7822 **CHANGE HISTORY**

7823 First released in Issue 1. Derived from Issue 1 of the SVID.

7824 **Issue 4**7825 The `<sys/types.h>` and `<sys/stat.h>` headers are now marked as optional (OH); these headers
7826 need not be included on XSI-conformant systems.

7827 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 7828
- The type of argument *path* is changed from `char*` to `const char*`.

7829 **Issue 6**7830 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.7831 The following new requirements on POSIX implementations derive from alignment with the
7832 Single UNIX Specification:

- 7833
- The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was
7834 required for conforming implementations of previous POSIX specifications, it was not
7835 required for UNIX applications.

7836 **NAME**7837 crypt — string encoding function (**CRYPT**)7838 **SYNOPSIS**7839 XSI `#include <unistd.h>`7840 `char *crypt(const char *key, const char *salt);`

7841

7842 **DESCRIPTION**7843 The *crypt()* function is a string encoding function. The algorithm is implementation-defined.7844 The *key* argument points to a string to be encoded. The *salt* argument is a string chosen from the
7845 set:

7846 a b c d e f g h i j k l m n o p q r s t u v w x y z

7847 A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

7848 0 1 2 3 4 5 6 7 8 9 . /

7849 The first two characters of this string may be used to perturb the encoding algorithm.

7850 The return value of *crypt()* points to static data that is overwritten by each call.7851 The *crypt()* function need not be reentrant. A function that is not required to be reentrant is not
7852 required to be thread-safe.7853 **RETURN VALUE**7854 Upon successful completion, *crypt()* shall return a pointer to the encoded string. The first two
7855 characters of the returned value are those of the *salt* argument. Otherwise, it shall return a null
7856 pointer and set *errno* to indicate the error.7857 **ERRORS**7858 The *crypt()* function shall fail if:

7859 [ENOSYS] The functionality is not supported on this implementation.

7860 **EXAMPLES**7861 **Encoding Passwords**7862 The following example finds a user database entry matching a particular user name and changes
7863 the current password to a new password. The *crypt()* function is used to generate an encoded
7864 version of each password. The first call to *crypt()* produces an encoded version of the old
7865 password; that encoded password is then compared to the password stored in the user database.
7866 The second call to *crypt()* encodes the new password before it is stored.7867 The *putpwent()* function, used in the following example, is not part of IEEE Std. 1003.1-200x.7868 `#include <unistd.h>`7869 `#include <pwd.h>`7870 `#include <string.h>`7871 `#include <stdio.h>`7872 `...`7873 `int valid_change;`7874 `int pfd; /* Integer for file descriptor returned by open(). */`7875 `FILE *fpfd; /* File pointer for use in putpwent(). */`7876 `struct passwd *p;`7877 `char user[100];`7878 `char oldpasswd[100];`7879 `char newpasswd[100];`

```

7880     char savepasswd[100];
7881     ...
7882     valid_change = 0;
7883     while ((p = getpwent()) != NULL) {
7884         /* Change entry if found. */
7885         if (strcmp(p->pw_name, user) == 0) {
7886             if (strcmp(p->pw_passwd, crypt(oldpasswd, p->pw_passwd)) == 0) {
7887                 strcpy(savepasswd, crypt(newpasswd, user));
7888                 p->pw_passwd = savepasswd;
7889                 valid_change = 1;
7890             }
7891             else {
7892                 fprintf(stderr, "Old password is not valid\n");
7893             }
7894         }
7895         /* Put passwd entry into ptmp. */
7896         putpwent(p, fpfd);
7897     }

```

7898 APPLICATION USAGE

7899 The values returned by this function need not be portable among XSI-conformant systems.

7900 RATIONALE

7901 None.

7902 FUTURE DIRECTIONS

7903 None.

7904 SEE ALSO

7905 *encrypt()*, *setkey()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

7906 CHANGE HISTORY

7907 First released in Issue 1. Derived from Issue 1 of the SVID.

7908 Issue 4

7909 The <**unistd.h**> header is added to the SYNOPSIS section.

7910 The type of arguments *key* and *salt* are changed from **char*** to **const char***.

7911 The DESCRIPTION now explicitly defines the characters that can appear in the *salt* argument.

7912 Issue 5

7913 Normative text previously in the APPLICATION USAGE section is moved to the
7914 DESCRIPTION.

7915 **NAME**

7916 csin, csinf, csinl — complex sine functions

7917 **SYNOPSIS**

7918 #include <complex.h>

7919 double complex csin(double complex *z*);7920 float complex csinf(float complex *z*);7921 long double complex csinl(long double complex *z*);7922 **DESCRIPTION**7923 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
7924 conflict between the requirements described here and the ISO C standard is unintentional. This
7925 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.7926 These functions shall compute the complex sine of *z*.7927 **RETURN VALUE**

7928 These functions shall return the complex sine value.

7929 **ERRORS**

7930 No errors are defined.

7931 **EXAMPLES**

7932 None.

7933 **APPLICATION USAGE**

7934 None.

7935 **RATIONALE**

7936 None.

7937 **FUTURE DIRECTIONS**

7938 None.

7939 **SEE ALSO**7940 *casin()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>7941 **CHANGE HISTORY**

7942 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7943 **NAME**

7944 csinh, csinhf, csinhl — complex hyperbolic sine functions

7945 **SYNOPSIS**

7946 #include <complex.h>

7947 double complex csinh(double complex *z*);7948 float complex csinhf(float complex *z*);7949 long double complex csinhl(long double complex *z*);7950 **DESCRIPTION**7951 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
7952 conflict between the requirements described here and the ISO C standard is unintentional. This
7953 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.7954 These functions shall compute the complex hyperbolic sine of *z*.7955 **RETURN VALUE**

7956 These functions shall return the complex hyperbolic sine value.

7957 **ERRORS**

7958 No errors are defined.

7959 **EXAMPLES**

7960 None.

7961 **APPLICATION USAGE**

7962 None.

7963 **RATIONALE**

7964 None.

7965 **FUTURE DIRECTIONS**

7966 None.

7967 **SEE ALSO**7968 *casinh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>7969 **CHANGE HISTORY**

7970 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

7971 **NAME**

7972 csqrt, csqrtf, csqrtl — complex square root functions

7973 **SYNOPSIS**

7974 #include <complex.h>

7975 double complex csqrt(double complex z);

7976 float complex csqrtf(float complex z);

7977 long double complex csqrtl(long double complex z);

7978 **DESCRIPTION**

7979 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
7980 conflict between the requirements described here and the ISO C standard is unintentional. This
7981 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

7982 These functions shall compute the complex square root of *z*, with a branch cut along the
7983 negative real axis.

7984 **RETURN VALUE**

7985 These functions shall return the complex square root value, in the range of the right half-plane
7986 (including the imaginary axis).

7987 **ERRORS**

7988 No errors are defined.

7989 **EXAMPLES**

7990 None.

7991 **APPLICATION USAGE**

7992 None.

7993 **RATIONALE**

7994 None.

7995 **FUTURE DIRECTIONS**

7996 None.

7997 **SEE ALSO**7998 *cabs()*, *cpow()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>7999 **CHANGE HISTORY**

8000 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

8001 **NAME**

8002 ctan, ctanf, ctanl — complex tangent functions

8003 **SYNOPSIS**

8004 #include <complex.h>

8005 double complex ctan(double complex *z*);8006 float complex ctanf(float complex *z*);8007 long double complex ctanl(long double complex *z*);8008 **DESCRIPTION**8009 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
8010 conflict between the requirements described here and the ISO C standard is unintentional. This
8011 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.8012 These functions shall compute the complex tangent of *z*.8013 **RETURN VALUE**

8014 These functions shall return the complex tangent value.

8015 **ERRORS**

8016 No errors are defined.

8017 **EXAMPLES**

8018 None.

8019 **APPLICATION USAGE**

8020 None.

8021 **RATIONALE**

8022 None.

8023 **FUTURE DIRECTIONS**

8024 None.

8025 **SEE ALSO**8026 *catan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**complex.h**>8027 **CHANGE HISTORY**

8028 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

8029 **NAME**

8030 ctanh, ctanhf, ctanhl — complex hyperbolic tangent functions

8031 **SYNOPSIS**

8032 #include <complex.h>

8033 double complex ctanh(double complex *z*);8034 float complex ctanhf(float complex *z*);8035 long double complex ctanhl(long double complex *z*);8036 **DESCRIPTION**8037 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
8038 conflict between the requirements described here and the ISO C standard is unintentional. This
8039 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.8040 These functions shall compute the complex hyperbolic tangent of *z*.8041 **RETURN VALUE**

8042 These functions shall return the complex hyperbolic tangent value.

8043 **ERRORS**

8044 No errors are defined.

8045 **EXAMPLES**

8046 None.

8047 **APPLICATION USAGE**

8048 None.

8049 **RATIONALE**

8050 None.

8051 **FUTURE DIRECTIONS**

8052 None.

8053 **SEE ALSO**8054 *catanh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <complex.h>8055 **CHANGE HISTORY**

8056 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

8057 **NAME**

8058 ctermid — generate a path name for controlling terminal

8059 **SYNOPSIS**

8060 #include <stdio.h>

8061 char *ctermid(char *s);

8062 **DESCRIPTION**

8063 The *ctermid()* function shall generate a string that, when used as a path name, refers to the
8064 current controlling terminal for the current process. If *ctermid()* returns a path name, access to
8065 the file is not guaranteed.

8066 If the application uses any of the `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS`
8067 functions, it shall ensure that the *ctermid()* function is called with a non-NULL parameter.

8068 **RETURN VALUE**

8069 If *s* is a null pointer, the string is generated in an area that may be static (and therefore may be
8070 overwritten by each call), the address of which shall be returned. Otherwise, *s* is assumed to
8071 point to a character array of at least `{L_ctermid}` bytes; the string is placed in this array and the
8072 value of *s* shall be returned. The symbolic constant `{L_ctermid}` is defined in `<stdio.h>`, and shall
8073 have a value greater than 0.

8074 The *ctermid()* function shall return an empty string if the path name that would refer to the
8075 controlling terminal cannot be determined, or if the function is unsuccessful.

8076 **ERRORS**

8077 No errors are defined.

8078 **EXAMPLES**8079 **Determining the Controlling Terminal for the Current Process**

8080 The following example returns a pointer to a string that identifies the controlling terminal for the
8081 current process. The path name for the terminal is stored in the array pointed to by the *ptr*
8082 argument, which has a size of `{L_ctermid}` bytes, as indicated by the *term* argument.

8083 #include <stdio.h>

8084 ...

8085 char term[L_ctermid];

8086 char *ptr;

8087 ptr = ctermid(term);

8088 **APPLICATION USAGE**

8089 The difference between *ctermid()* and *ttyname()* is that *ttyname()* must be handed a file
8090 descriptor and return a path of the terminal associated with that file descriptor, while *ctermid()*
8091 returns a string (such as `"/dev/tty"`) that refers to the current controlling terminal if used as a
8092 path name.

8093 **RATIONALE**

8094 `{L_ctermid}` must be defined appropriately for a given implementation and must be greater than
8095 zero so that array declarations using it are accepted by the compiler. The value includes the
8096 terminating null byte.

8097 Portable applications that use threads cannot call *ctermid()* with NULL as the parameter if either
8098 `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS` is defined. If *s* is not NULL, the
8099 *ctermid()* function generates a string that, when used as a path name, refers to the current
8100 controlling terminal for the current process. If *s* is NULL, the return value of *ctermid()* is

- 8101 undefined.
- 8102 If the *ctermid()* function returns a path name, access to the file is not guaranteed.
- 8103 There is no additional burden on the programmer—changing to use a hypothetical thread-safe
8104 version of *ctermid()* along with allocating a buffer is more of a burden than merely allocating a
8105 buffer. Application code should not assume that the returned string is short, as some
8106 implementations have more than two path name components before reaching a logical device
8107 name.
- 8108 **FUTURE DIRECTIONS**
- 8109 None.
- 8110 **SEE ALSO**
- 8111 *ttynname()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>
- 8112 **CHANGE HISTORY**
- 8113 First released in Issue 1. Derived from Issue 1 of the SVID.
- 8114 **Issue 4**
- 8115 The following change is incorporated for alignment with the ISO POSIX-1 standard:
- 8116 • The DESCRIPTION and RETURN VALUE sections, though functionally identical to Issue 3,
8117 are rewritten.
- 8118 **Issue 5**
- 8119 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.
- 8120 **Issue 6**
- 8121 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

8122 **NAME**

8123 ctime, ctime_r — convert a time value to date and time string

8124 **SYNOPSIS**

8125 #include <time.h>

8126 char *ctime(const time_t *clock);

8127 TSF char *ctime_r(const time_t *clock, char *buf);

8128

8129 **DESCRIPTION**

8130 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 8131 conflict between the requirements described here and the ISO C standard is unintentional. This
 8132 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

8133 The *ctime()* function shall convert the time pointed to by *clock*, representing time in seconds
 8134 since the Epoch, to local time in the form of a string. It is equivalent to:

8135 asctime(localtime(clock))

8136 CX The *asctime()*, *ctime()*, *gmtime()*, and *localtime()* functions return values in one of two static
 8137 objects: a broken-down time structure and an array of **char**. Execution of any of the functions
 8138 may overwrite the information returned in either of these objects by any of the other functions.

8139 The *ctime()* function need not be reentrant. A function that is not required to be reentrant is not
 8140 required to be thread-safe.

8141 TSF The *ctime_r()* function shall convert the calendar time pointed to by *clock* to local time in exactly
 8142 the same form as *ctime()* and puts the string into the array pointed to by *buf* (which contains at
 8143 least 26 bytes) and return *buf*.

8144 Unlike *ctime()*, the thread-safe version *ctime_r()* is not required to set *tzname*.

8145 **RETURN VALUE**

8146 The *ctime()* function shall return the pointer returned by *asctime()* with that broken-down time
 8147 as an argument.

8148 TSF Upon successful completion, *ctime_r()* shall return a pointer to the string pointed to by *buf*.
 8149 When an error is encountered, a null pointer shall be returned.

8150 **ERRORS**

8151 No errors are defined.

8152 **EXAMPLES**

8153 None.

8154 **APPLICATION USAGE**

8155 Values for the broken-down time structure can be obtained by calling *gmtime()* or *localtime()*.
 8156 The *ctime()* function is included for compatibility with older implementations, and does not
 8157 support localized date and time formats. Applications should use the *strftime()* function to
 8158 achieve maximum portability.

8159 The *ctime_r()* function is thread-safe and shall return values in a user-supplied buffer instead of
 8160 possibly using a static data area that may be overwritten by each call.

8161 **RATIONALE**

8162 None.

8163 **FUTURE DIRECTIONS**

8164 None.

8165 **SEE ALSO**8166 *asctime()*, *clock()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strftime()*, *strptime()*, *time()*, *utime()*,
8167 the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>8168 **CHANGE HISTORY**

8169 First released in Issue 1. Derived from Issue 1 of the SVID.

8170 **Issue 4**8171 The APPLICATION USAGE section is expanded to describe the time-handling functions
8172 generally and to refer users to *strftime()*, which is a locale-dependent time-handling function.

8173 The following change is incorporated for alignment with the ISO C standard:

- 8174
- The type of argument *clock* is changed from **time_t*** to **const time_t***.

8175 **Issue 5**8176 Normative text previously in the APPLICATION USAGE section is moved to the
8177 DESCRIPTION.8178 The *ctime_r()* function is included for alignment with the POSIX Threads Extension.8179 A note indicating that the *ctime()* function need not be reentrant is added to the DESCRIPTION.8180 **Issue 6**

8181 Extensions beyond the ISO C standard are now marked.

8182 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

8183 The APPLICATION USAGE section is updated to include a note on the thread-safe function and
8184 its avoidance of possibly using a static data area.

8185 **NAME**

8186 daylight — daylight savings time flag

8187 **SYNOPSIS**

8188 xSI #include <time.h>

8189 extern int daylight;

8190

8191 **DESCRIPTION**

8192 Refer to *tzset()*.

8193 NAME

8194 dbm_clearerr, dbm_close, dbm_delete, dbm_error, dbm_fetch, dbm_firstkey, dbm_nextkey,
8195 dbm_open, dbm_store — database functions

8196 SYNOPSIS

```
8197 xSI #include <ndbm.h>

8198 int dbm_clearerr(DBM *db);
8199 void dbm_close(DBM *db);
8200 int dbm_delete(DBM *db, datum key);
8201 int dbm_error(DBM *db);
8202 datum dbm_fetch(DBM *db, datum key);
8203 datum dbm_firstkey(DBM *db);
8204 datum dbm_nextkey(DBM *db);
8205 DBM *dbm_open(const char *file, int open_flags, mode_t file_mode);
8206 int dbm_store(DBM *db, datum key, datum content, int store_mode);
8207
```

8208 DESCRIPTION

8209 These functions create, access, and modify a database.

8210 A **datum** consists of at least two members, *dptr* and *dsize*. The *dptr* member points to an object
8211 that is *dsize* bytes in length. Arbitrary binary data, as well as character strings, may be stored in
8212 the object pointed to by *dptr*.

8213 The database is stored in two files. One file is a directory containing a bit map of keys and has
8214 **.dir** as its suffix. The second file contains all data and has **.pag** as its suffix.

8215 The *dbm_open()* function shall open a database. The *file* argument to the function is the path
8216 name of the database. The function opens two files named *file.dir* and *file.pag*. The *open_flags*
8217 argument has the same meaning as the *flags* argument of *open()* except that a database opened
8218 for write-only access opens the files for read and write access and the behavior of the
8219 O_APPEND flag is unspecified. The *file_mode* argument has the same meaning as the third
8220 argument of *open()*.

8221 The *dbm_close()* function shall close a database. The application shall ensure that argument *db* is
8222 a pointer to a **dbm** structure that has been returned from a call to *dbm_open()*.

8223 The *dbm_fetch()* function shall read a record from a database. The argument *db* is a pointer to a
8224 database structure that has been returned from a call to *dbm_open()*. The argument *key* is a
8225 **datum** that has been initialized by the application to the value of the key that matches the key of
8226 the record the program is fetching.

8227 The *dbm_store()* function shall write a record to a database. The argument *db* is a pointer to a
8228 database structure that has been returned from a call to *dbm_open()*. The argument *key* is a
8229 **datum** that has been initialized by the application to the value of the key that identifies (for
8230 subsequent reading, writing, or deleting) the record the application is writing. The argument
8231 *content* is a **datum** that has been initialized by the application to the value of the record the
8232 program is writing. The argument *store_mode* controls whether *dbm_store()* replaces any pre-
8233 existing record that has the same key that is specified by the *key* argument. The application shall
8234 set *store_mode* to either DBM_INSERT or DBM_REPLACE. If the database contains a record that
8235 matches the *key* argument and *store_mode* is DBM_REPLACE, the existing record is replaced with
8236 the new record. If the database contains a record that matches the *key* argument and *store_mode*
8237 is DBM_INSERT, the existing record is left unchanged and the new record ignored. If the
8238 database does not contain a record that matches the *key* argument and *store_mode* is either
8239 DBM_INSERT or DBM_REPLACE, the new record is inserted in the database.

8240 The application shall ensure that the sum of the sizes of a key/content pair does not exceed the
8241 internal block size. Moreover, the application shall ensure that all key/content pairs that hash
8242 together fit on a single block. The *dbm_store()* function shall return an error in the event that a
8243 disk block fills with inseparable data.

8244 The *dbm_delete()* function shall delete a record and its key from the database. The argument *db* is
8245 a pointer to a database structure that has been returned from a call to *dbm_open()*. The argument
8246 *key* is a **datum** that has been initialized by the application to the value of the key that identifies
8247 the record the program is deleting.

8248 The *dbm_firstkey()* function shall return the first key in the database. The argument *db* is a
8249 pointer to a database structure that has been returned from a call to *dbm_open()*.

8250 The *dbm_nextkey()* function shall return the next key in the database. The argument *db* is a
8251 pointer to a database structure that has been returned from a call to *dbm_open()*. The application
8252 shall ensure that the *dbm_firstkey()* function is called before calling *dbm_nextkey()*. Subsequent
8253 calls to *dbm_nextkey()* return the next key until all of the keys in the database have been
8254 returned.

8255 The *dbm_error()* function shall return the error condition of the database. The argument *db* is a
8256 pointer to a database structure that has been returned from a call to *dbm_open()*.

8257 The *dbm_clearerr()* function shall clear the error condition of the database. The argument *db* is a
8258 pointer to a database structure that has been returned from a call to *dbm_open()*.

8259 These database functions shall support an internal block size large enough to support
8260 key/content pairs of at least 1 023 bytes.

8261 The *dptr* pointers returned by these functions may point into static storage that may be changed
8262 by subsequent calls.

8263 These functions need not be reentrant. A function that is not required to be reentrant is not
8264 required to be thread-safe.

8265 RETURN VALUE

8266 The *dbm_store()* and *dbm_delete()* functions shall return 0 when they succeed and a negative
8267 value when they fail.

8268 The *dbm_store()* function shall return 1 if it is called with a *flags* value of DBM_INSERT and the
8269 function finds an existing record with the same key.

8270 The *dbm_error()* function shall return 0 if the error condition is not set and return a non-zero
8271 value if the error condition is set.

8272 The return value of *dbm_clearerr()* is unspecified.

8273 The *dbm_firstkey()* and *dbm_nextkey()* functions shall return a key **datum**. When the end of the
8274 database is reached, the *dptr* member of the key is a null pointer. If an error is detected, the *dptr*
8275 member of the key shall be a null pointer and the error condition of the database shall be set.

8276 The *dbm_fetch()* function shall return a content **datum**. If no record in the database matches the
8277 key or if an error condition has been detected in the database, the *dptr* member of the content
8278 shall be a null pointer.

8279 The *dbm_open()* function shall return a pointer to a database structure. If an error is detected
8280 during the operation, *dbm_open()* shall return a **(DBM*)0**.

8281 ERRORS

8282 No errors are defined.

8283 EXAMPLES

8284 None.

8285 APPLICATION USAGE

8286 The following code can be used to traverse the database:

```
8287 for(key = dbm_firstkey(db); key.dptr != NULL; key = dbm_nextkey(db))
```

8288 The *dbm_* functions provided in this library should not be confused in any way with those of a
8289 general-purpose database management system. These functions do not provide for multiple
8290 search keys per entry, they do not protect against multi-user access (in other words they do not
8291 lock records or files), and they do not provide the many other useful database functions that are
8292 found in more robust database management systems. Creating and updating databases by use of
8293 these functions is relatively slow because of data copies that occur upon hash collisions. These
8294 functions are useful for applications requiring fast lookup of relatively static information that is
8295 to be indexed by a single key.

8296 The *dbm_delete()* function need not physically reclaim file space, although it does make it
8297 available for reuse by the database.

8298 After calling *dbm_store()* or *dbm_delete()* during a pass through the keys by *dbm_firstkey()* and
8299 *dbm_nextkey()*, the application should reset the database by calling *dbm_firstkey()* before again
8300 calling *dbm_nextkey()*. The contents of these files are unspecified and may not be portable.

8301 RATIONALE

8302 None.

8303 FUTURE DIRECTIONS

8304 None.

8305 SEE ALSO

8306 *open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**ndbm.h**>

8307 CHANGE HISTORY

8308 First released in Issue 4, Version 2.

8309 Issue 5

8310 Moved from X/OPEN UNIX extension to BASE.

8311 Normative text previously in the APPLICATION USAGE section is moved to the
8312 DESCRIPTION.

8313 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

8314 Issue 6

8315 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

8316 **NAME**

8317 difftime — compute the difference between two calendar time values

8318 **SYNOPSIS**

8319 #include <time.h>

8320 double difftime(time_t *time1*, time_t *time0*);

8321 **DESCRIPTION**

8322 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
8323 conflict between the requirements described here and the ISO C standard is unintentional. This
8324 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

8325 The *difftime()* function shall compute the difference between two calendar times (as returned by
8326 *time()*): *time1*– *time0*.

8327 **RETURN VALUE**

8328 The *difftime()* function shall return the difference expressed in seconds as a type **double**.

8329 **ERRORS**

8330 No errors are defined.

8331 **EXAMPLES**

8332 None.

8333 **APPLICATION USAGE**

8334 None.

8335 **RATIONALE**

8336 None.

8337 **FUTURE DIRECTIONS**

8338 None.

8339 **SEE ALSO**

8340 *asctime()*, *clock()*, *ctime()*, *gmtime()*, *localtime()*, *mktime()*, *strftime()*, *strptime()*, *time()*, *utime()*,
8341 the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

8342 **CHANGE HISTORY**

8343 First released in Issue 4. Derived from the ISO C standard.

8344 **NAME**

8345 dirname — report the parent directory name of a file path name

8346 **SYNOPSIS**

8347 xSI #include <libgen.h>

8348 char *dirname(char *path);

8349

8350 **DESCRIPTION**8351 The *dirname()* function shall take a pointer to a character string that contains a path name, and
8352 return a pointer to a string that is a path name of the parent directory of that file. Trailing '/'
8353 characters in the path are not counted as part of the path.8354 If *path* does not contain a '/', then *dirname()* shall return a pointer to the string ".". If *path* is a
8355 null pointer or points to an empty string, *dirname()* shall return a pointer to the string ".".8356 The *dirname()* function need not be reentrant. A function that is not required to be reentrant is
8357 not required to be thread-safe.8358 **RETURN VALUE**8359 The *dirname()* function shall return a pointer to a string that is the parent directory of *path*. If
8360 *path* is a null pointer or points to an empty string, a pointer to a string "." is returned.8361 The *dirname()* function may modify the string pointed to by *path*, and may return a pointer to
8362 static storage that may then be overwritten by subsequent calls to *dirname()*.8363 **ERRORS**

8364 No errors are defined.

8365 **EXAMPLES**8366 The following code fragment reads a path name, changes the current working directory to the
8367 parent directory, and opens the file.8368 char path[MAXPATHLEN], *pathcopy;
8369 int fd;
8370 fgets(path, MAXPATHLEN, stdin);
8371 pathcopy = strdup(path);
8372 chdir(dirname(pathcopy));
8373 fd = open(basename(path), O_RDONLY);8374 **Sample Input and Output Strings for dirname()**8375 In the following table, the input string is the value pointed to by *path*, and the output string is
8376 the return value of the *dirname()* function.

8377

8378

8379

8380

8381

8382

8383

Input String	Output String
"/usr/lib"	"/usr"
"/usr/"	"/"
"usr"	."
"/"	"/"
."	."
.."	."

8384 **Changing the Current Directory to the Parent Directory**

8385 The following program fragment reads a path name, changes the current working directory to
8386 the parent directory, and opens the file.

```
8387 #include <unistd.h>
8388 #include <limits.h>
8389 #include <stdio.h>
8390 #include <fcntl.h>
8391 #include <string.h>
8392 #include <libgen.h>
8393 ...
8394 char path[PATH_MAX], *pathcopy;
8395 int fd;
8396 ...
8397 fgets(path, PATH_MAX, stdin);
8398 pathcopy = strdup(path);
8399 chdir(dirname(pathcopy));
8400 fd = open(basename(path), O_RDONLY);
```

8401 **APPLICATION USAGE**

8402 The *dirname()* and *basename()* functions together yield a complete path name. The expression
8403 *dirname(path)* obtains the path name of the directory where *basename(path)* is found.

8404 Since the meaning of the leading *"/"* is implementation-defined, *dirname("/foo)* may return
8405 either *"/"* or *'/'* (but nothing else).

8406 **RATIONALE**

8407 None.

8408 **FUTURE DIRECTIONS**

8409 None.

8410 **SEE ALSO**

8411 *basename()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**libgen.h**>

8412 **CHANGE HISTORY**

8413 First released in Issue 4, Version 2.

8414 **Issue 5**

8415 Moved from X/OPEN UNIX extension to BASE.

8416 Normative text previously in the APPLICATION USAGE section is moved to the
8417 DESCRIPTION.

8418 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

8419 **NAME**

8420 `div` — compute the quotient and remainder of an integer division

8421 **SYNOPSIS**

8422 `#include <stdlib.h>`

8423 `div_t div(int numer, int denom);`

8424 **DESCRIPTION**

8425 `CX` The functionality described on this reference page is aligned with the ISO C standard. Any
8426 conflict between the requirements described here and the ISO C standard is unintentional. This
8427 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

8428 The `div()` function shall compute the quotient and remainder of the division of the numerator
8429 *numer* by the denominator *denom*. If the division is inexact, the resulting quotient is the integer
8430 of lesser magnitude that is the nearest to the algebraic quotient. If the result cannot be
8431 represented, the behavior is undefined; otherwise, *quot*denom+rem* shall equal *numer*.

8432 **RETURN VALUE**

8433 The `div()` function shall return a structure of type `div_t`, comprising both the quotient and the
8434 remainder. The structure includes the following members, in any order:

8435 `int quot; /* quotient */`

8436 `int rem; /* remainder */`

8437 **ERRORS**

8438 No errors are defined.

8439 **EXAMPLES**

8440 None.

8441 **APPLICATION USAGE**

8442 None.

8443 **RATIONALE**

8444 None.

8445 **FUTURE DIRECTIONS**

8446 None.

8447 **SEE ALSO**

8448 `ldiv()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>`

8449 **CHANGE HISTORY**

8450 First released in Issue 4. Derived from the ISO C standard.

8451 **NAME**8452 dldclose — close a *dlopen()* object8453 **SYNOPSIS**

```
8454 xSI       #include <dldfcn.h>
8455           int dldclose(void *handle);
8456
```

8457 **DESCRIPTION**

8458 The *dldclose()* function is used to inform the system that the object referenced by a *handle* returned
8459 from a previous *dlopen()* invocation is no longer needed by the application.

8460 The use of *dldclose()* reflects a statement of intent on the part of the process, but does not create
8461 any requirement upon the implementation, such as removal of the code or symbols referenced
8462 by *handle*. Once an object has been closed using *dldclose()* an application should assume that its
8463 symbols are no longer available to *dlsym()*. All objects loaded automatically as a result of
8464 invoking *dlopen()* on the referenced object are also closed if this is the last reference to it.

8465 Although a *dldclose()* operation is not required to remove structures from an address space,
8466 neither is an implementation prohibited from doing so. The only restriction on such a removal is
8467 that no object shall be removed to which references have been relocated, until or unless all such
8468 references are removed. For instance, an object that had been loaded with a *dlopen()* operation
8469 specifying the `RTLD_GLOBAL` flag might provide a target for dynamic relocations performed in
8470 the processing of other objects—in such environments, an application may assume that no
8471 relocation, once made, shall be undone or remade unless the object requiring the relocation has
8472 itself been removed.

8473 **RETURN VALUE**

8474 If the referenced object was successfully closed, *dldclose()* shall return 0. If the object could not be
8475 closed, or if *handle* does not refer to an open object, *dldclose()* shall return a non-zero value. More
8476 detailed diagnostic information shall be available through *dlderror()*.

8477 **ERRORS**

8478 No errors are defined.

8479 **EXAMPLES**8480 The following example illustrates use of *dlopen()* and *dldclose()*:

```
8481           ...
8482           /* Open a dynamic library and then close it ... */
8483           #include <dldfcn.h>
8484           void *mylib;
8485           int  eret;
8486           mylib = dlopen("mylib.so.1", RTLD_LAZY);
8487           ...
8488           eret = dldclose(mylib);
8489           ...
```

8490 **APPLICATION USAGE**

8491 A portable application should employ a *handle* returned from a *dlopen()* invocation only within a
8492 given scope bracketed by the *dlopen()* and *dldclose()* operations. Implementations are free to use
8493 reference counting or other techniques such that multiple calls to *dlopen()* referencing the same
8494 object may return the same object for *handle*. Implementations are also free to reuse a *handle*.
8495 For these reasons, the value of a *handle* must be treated as an opaque object by the application,
8496 used only in calls to *dlsym()* and *dldclose()*.

8497 **RATIONALE**

8498 None.

8499 **FUTURE DIRECTIONS**

8500 None.

8501 **SEE ALSO**8502 *derror()*, *dlopen()*, *dlsym()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**dlfcn.h**>8503 **CHANGE HISTORY**

8504 First released in Issue 5.

8505 **Issue 6**8506 The DESCRIPTION is updated to say that the referenced object is closed “if this is the last
8507 reference to it”.

8508 **NAME**8509 `dlderror` — get diagnostic information8510 **SYNOPSIS**8511 XSI `#include <dldfcn.h>`8512 `char *dlderror(void);`

8513

8514 **DESCRIPTION**

8515 The `dlderror()` function shall return a null-terminated character string (with no trailing <newline>)
8516 that describes the last error that occurred during dynamic linking processing. If no dynamic
8517 linking errors have occurred since the last invocation of `dlderror()`, `dlderror()` shall return NULL.
8518 Thus, invoking `dlderror()` a second time, immediately following a prior invocation, shall result in
8519 NULL being returned.

8520 The `dlderror()` function need not be reentrant. A function that is not required to be reentrant is not
8521 required to be thread-safe.

8522 **RETURN VALUE**

8523 If successful, `dlderror()` shall return a null-terminated character string; otherwise, NULL shall be
8524 returned.

8525 **ERRORS**

8526 No errors are defined.

8527 **EXAMPLES**

8528 The following example prints out the last dynamic linking error:

```
8529 ...  
8530 #include <dldfcn.h>  
8531 char *errstr;  
8532 errstr = dlderror();  
8533 if (errstr != NULL)  
8534 printf ("A dynamic linking error occurred: (%s)\n", errstr);  
8535 ...
```

8536 **APPLICATION USAGE**

8537 The messages returned by `dlderror()` may reside in a static buffer that is overwritten on each call
8538 to `dlderror()`. Application code should not write to this buffer. Programs wishing to preserve an
8539 error message should make their own copies of that message. Depending on the application
8540 environment with respect to asynchronous execution events, such as signals or other
8541 asynchronous computation sharing the address space, portable applications should use a critical
8542 section to retrieve the error pointer and buffer.

8543 **RATIONALE**

8544 None.

8545 **FUTURE DIRECTIONS**

8546 None.

8547 **SEE ALSO**8548 `dldclose()`, `dldopen()`, `dldsym()`, the Base Definitions volume of IEEE Std. 1003.1-200x, <dldfcn.h>

8549 **CHANGE HISTORY**

8550 First released in Issue 5.

8551 **Issue 6**

8552 In the DESCRIPTION the note about reentrancy and thread-safety is added.

8553 NAME

8554 dlopen — gain access to an executable object file

8555 SYNOPSIS

8556 xSI #include <dlfcn.h>

8557 void *dlopen(const char *file, int mode);

8558

8559 DESCRIPTION

8560 The *dlopen()* function shall make an executable object file specified by *file* available to the calling
 8561 program. The class of files eligible for this operation and the manner of their construction are
 8562 specified by the implementation, though typically such files are executable objects such as
 8563 shared libraries, relocatable files, or programs. Note that some implementations permit the
 8564 construction of dependencies between such objects that are embedded within files. In such
 8565 cases, a *dlopen()* operation shall load such dependencies in addition to the object referenced by
 8566 *file*. Implementations may also impose specific constraints on the construction of programs that
 8567 can employ *dlopen()* and its related services.

8568 A successful *dlopen()* shall return a *handle* which the caller may use on subsequent calls to
 8569 *dlsym()* and *dlclose()*. The value of this *handle* should not be interpreted in any way by the caller.

8570 *file* is used to construct a path name to the object file. If *file* contains a slash character, the *file*
 8571 argument is used as the path name for the file. Otherwise, *file* is used in an implementation-
 8572 defined manner to yield a path name.

8573 If the value of *file* is 0, *dlopen()* shall provide a *handle* on a global symbol object. This object
 8574 provides access to the symbols from an ordered set of objects consisting of the original program
 8575 image file, together with any objects loaded at program start-up as specified by that process
 8576 image file (for example, shared libraries), and the set of objects loaded using a *dlopen()* operation
 8577 together with the RTLD_GLOBAL flag. As the latter set of objects can change during execution,
 8578 the set identified by *handle* can also change dynamically.

8579 Only a single copy of an object file is brought into the address space, even if *dlopen()* is invoked
 8580 multiple times in reference to the file, and even if different path names are used to reference the
 8581 file.

8582 The *mode* parameter describes how *dlopen()* shall operate upon *file* with respect to the processing
 8583 of relocations and the scope of visibility of the symbols provided within *file*. When an object is
 8584 brought into the address space of a process, it may contain references to symbols whose
 8585 addresses are not known until the object is loaded. These references shall be relocated before the
 8586 symbols can be accessed. The *mode* parameter governs when these relocations take place and
 8587 may have the following values:

8588 RTLD_LAZY Relocations shall be performed at an implementation-defined time,
 8589 ranging from the time of the *dlopen()* call until the first reference to a
 8590 given symbol occurs. Specifying RTLD_LAZY should improve
 8591 performance on implementations supporting dynamic symbol binding as
 8592 a process may not reference all of the functions in any given object. And,
 8593 for systems supporting dynamic symbol resolution for normal process
 8594 execution, this behavior mimics the normal handling of process
 8595 execution.

8596 RTLD_NOW All necessary relocations shall be performed when the object is first
 8597 loaded. This may waste some processing if relocations are performed for
 8598 functions that are never referenced. This behavior may be useful for
 8599 applications that need to know as soon as an object is loaded that all

8600 symbols referenced during execution are available.

8601 Any object loaded by *dlopen()* that requires relocations against global symbols can reference the
8602 symbols in the original process image file, any objects loaded at program start-up, from the
8603 object itself as well as any other object included in the same *dlopen()* invocation, and any objects
8604 that were loaded in any *dlopen()* invocation and which specified the RTLD_GLOBAL flag. To
8605 determine the scope of visibility for the symbols loaded with a *dlopen()* invocation, the *mode*
8606 parameter should be a bitwise-inclusive OR with one of the following values:

8607 RTLD_GLOBAL The object's symbols shall be made available for the relocation processing
8608 of any other object. In addition, symbol lookup using *dlopen(0, mode)* and
8609 an associated *dlsym()* allows objects loaded with this *mode* to be searched.

8610 RTLD_LOCAL The object's symbols shall not be made available for the relocation
8611 processing of any other object.

8612 If neither RTLD_GLOBAL nor RTLD_LOCAL are specified, then an implementation-defined
8613 default behavior shall be applied.

8614 If a *file* is specified in multiple *dlopen()* invocations, *mode* is interpreted at each invocation. Note,
8615 however, that once RTLD_NOW has been specified all relocations shall have been completed
8616 rendering further RTLD_NOW operations redundant and any further RTLD_LAZY operations
8617 irrelevant. Similarly, note that once RTLD_GLOBAL has been specified the object shall maintain
8618 the RTLD_GLOBAL status regardless of any previous or future specification of RTLD_LOCAL,
8619 as long as the object remains in the address space (see *dlclose()*).

8620 Symbols introduced into a program through calls to *dlopen()* may be used in relocation
8621 activities. Symbols so introduced may duplicate symbols already defined by the program or
8622 previous *dlopen()* operations. To resolve the ambiguities such a situation might present, the
8623 resolution of a symbol reference to symbol definition is based on a symbol resolution order. Two
8624 such resolution orders are defined: *load* or *dependency* ordering. Load order establishes an
8625 ordering among symbol definitions, such that the definition first loaded (including definitions
8626 from the image file and any dependent objects loaded with it) has priority over objects added
8627 later (via *dlopen()*). Load ordering is used in relocation processing. Dependency ordering uses a
8628 breadth-first order starting with a given object, then all of its dependencies, then any dependents
8629 of those, iterating until all dependencies are satisfied. With the exception of the global symbol
8630 object obtained via a *dlopen()* operation on a *file* of 0, dependency ordering is used by the
8631 *dlsym()* function. Load ordering is used in *dlsym()* operations upon the global symbol object.

8632 When an object is first made accessible via *dlopen()* it and its dependent objects are added in
8633 dependency order. Once all the objects are added, relocations are performed using load order.
8634 Note that if an object or its dependencies had been previously loaded, the load and dependency
8635 orders may yield different resolutions.

8636 The symbols introduced by *dlopen()* operations, and available through *dlsym()* are at a
8637 minimum those which are exported as symbols of global scope by the object. Typically such
8638 symbols shall be those that were specified in (for example) C source code as having *extern*
8639 linkage. The precise manner in which an implementation constructs the set of exported symbols
8640 for a *dlopen()* object is specified by that implementation.

8641 **RETURN VALUE**

8642 If *file* cannot be found, cannot be opened for reading, is not of an appropriate object format for
8643 processing by *dlopen()*, or if an error occurs during the process of loading *file* or relocating its
8644 symbolic references, *dlopen()* shall return NULL. More detailed diagnostic information shall be
8645 available through *dlerror()*.

8646 **ERRORS**

8647 No errors are defined.

8648 **EXAMPLES**

8649 None.

8650 **APPLICATION USAGE**

8651 None.

8652 **RATIONALE**

8653 None.

8654 **FUTURE DIRECTIONS**

8655 None.

8656 **SEE ALSO**

8657 *dlclose()*, *dLError()*, *dlsym()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**dlfcn.h**>

8658 **CHANGE HISTORY**

8659 First released in Issue 5.

8660 **NAME**8661 dlsym — obtain the address of a symbol from a *dlopen()* object8662 **SYNOPSIS**

8663 XSI #include <dlfcn.h>

8664 void *dlsym(void *restrict *handle*, const char *restrict *name*);

8665

8666 **DESCRIPTION**8667 The *dlsym()* function allows a process to obtain the address of a symbol defined within an object
8668 made accessible through a *dlopen()* call. *handle* is the value returned from a call to *dlopen()* (and
8669 which has not since been released via a call to *dlclose()*), and *name* is the symbol's name as a
8670 character string.8671 The *dlsym()* function shall search for the named symbol in all objects loaded automatically as a
8672 result of loading the object referenced by *handle* (see *dlopen()*). Load ordering is used in *dlsym()*
8673 operations upon the global symbol object. The symbol resolution algorithm used shall be
8674 dependency order as described in *dlopen()*.

8675 The RTLD_NEXT flag is reserved for future use.

8676 **RETURN VALUE**8677 If *handle* does not refer to a valid object opened by *dlopen()*, or if the named symbol cannot be
8678 found within any of the objects associated with *handle*, *dlsym()* shall return NULL. More
8679 detailed diagnostic information shall be available through *dLError()*.8680 **ERRORS**

8681 No errors are defined.

8682 **EXAMPLES**8683 The following example shows how *dlopen()* and *dlsym()* can be used to access either function or
8684 data objects. For simplicity, error checking has been omitted.8685 void *handle;
8686 int *iptr, (*fptr)(int);

8687 /* open the needed object */
8688 handle = dlopen("/usr/home/me/libfoo.so.1", RTLD_LAZY);

8689 /* find the address of function and data objects */
8690 fptr = (int (*)(int))dlsym(handle, "my_function");
8691 iptr = (int *)dlsym(handle, "my_object");

8692 /* invoke function, passing value of integer as a parameter */
8693 (*fptr)(*iptr);8694 **APPLICATION USAGE**8695 Special purpose values for *handle* are reserved for future use. These values and their meanings
8696 are:8697 RTLD_NEXT Specifies the next object after this one that defines *name*. *This one* refers to the
8698 object containing the invocation of *dlsym()*. The *next* object is the one found
8699 upon the application of a load order symbol resolution algorithm (see
8700 *dlopen()*). The next object is either one of global scope (because it was
8701 introduced as part of the original process image or because it was added with
8702 a *dlopen()* operation including the RTLD_GLOBAL flag), or is an object that
8703 was included in the same *dlopen()* operation that loaded this one.

8704 The RTLD_NEXT flag is useful to navigate an intentionally created hierarchy
8705 of multiply-defined symbols created through *interposition*. For example, if a
8706 program wished to create an implementation of *malloc()* that embedded some
8707 statistics gathering about memory allocations, such an implementation could
8708 use the real *malloc()* definition to perform the memory allocation—and itself
8709 only embed the necessary logic to implement the statistics gathering function.

8710 **RATIONALE**

8711 None.

8712 **FUTURE DIRECTIONS**

8713 None.

8714 **SEE ALSO**

8715 *dlclose()*, *derror()*, *dlopen()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <dlfcn.h>

8716 **CHANGE HISTORY**

8717 First released in Issue 5.

8718 **Issue 6**

8719 The **restrict** keyword is added to the *dlsym()* prototype for alignment with the
8720 ISO/IEC 9899:1999 standard.

8721 NAME

8722 drand48, erand48, jrand48, lcong48, lrand48, mrand48, nrand48, seed48, srand48 — generate
8723 uniformly distributed pseudo-random numbers

8724 SYNOPSIS

```
8725 xSI #include <stdlib.h>

8726 double drand48(void);
8727 double erand48(unsigned short xsubi[3]);
8728 long jrand48(unsigned short xsubi[3]);
8729 void lcong48(unsigned short param[7]);
8730 long lrand48(void);
8731 long mrand48(void);
8732 long nrand48(unsigned short xsubi[3]);
8733 unsigned short *seed48(unsigned short seed16v[3]);
8734 void srand48(long seedval);
8735
```

8736 DESCRIPTION

8737 This family of functions generates pseudo-random numbers using a linear congruential
8738 algorithm and 48-bit integer arithmetic.

8739 The *drand48()* and *erand48()* functions shall return non-negative, double-precision, floating-
8740 point values, uniformly distributed over the interval [0.0,1.0).

8741 The *lrnd48()* and *nrnd48()* functions shall return non-negative, long integers, uniformly
8742 distributed over the interval $[0, 2^{31})$.

8743 The *mrnd48()* and *jrnd48()* functions shall return signed long integers uniformly distributed
8744 over the interval $[-2^{31}, 2^{31})$.

8745 The *srand48()*, *seed48()*, and *lcong48()* are initialization entry points, one of which should be
8746 invoked before either *drand48()*, *lrnd48()*, or *mrnd48()* is called. (Although it is not
8747 recommended practice, constant default initializer values shall be supplied automatically if
8748 *drand48()*, *lrnd48()*, or *mrnd48()* is called without a prior call to an initialization entry point.)
8749 The *erand48()*, *nrnd48()*, and *jrnd48()* functions do not require an initialization entry point to
8750 be called first.

8751 All the routines work by generating a sequence of 48-bit integer values, X_i , according to the
8752 linear congruential formula:

$$8753 \quad X_{n+1} = (aX_n + c)_{\text{mod } m} \quad n \geq 0$$

8754 The parameter $m = 2^{48}$; hence 48-bit integer arithmetic is performed. Unless *lcong48()* is invoked,
8755 the multiplier value a and the addend value c are given by:

$$8756 \quad a = 5\text{DEECE66D}_{16} = 273673163155_8$$

$$8757 \quad c = \text{B}_{16} = 13_8$$

8758 The value returned by any of the *drand48()*, *erand48()*, *jrnd48()*, *lrnd48()*, *mrnd48()*, or
8759 *nrnd48()* functions is computed by first generating the next 48-bit X_i in the sequence. Then the
8760 appropriate number of bits, according to the type of data item to be returned, are copied from
8761 the high-order (leftmost) bits of X_i and transformed into the returned value.

8762 The *drand48()*, *lrnd48()*, and *mrnd48()* functions store the last 48-bit X_i generated in an
8763 internal buffer; that is why the application shall ensure that these are initialized prior to being
8764 invoked. The *erand48()*, *nrnd48()*, and *jrnd48()* functions require the calling program to
8765 provide storage for the successive X_i values in the array specified as an argument when the

8766 functions are invoked. That is why these routines do not have to be initialized; the calling
 8767 program merely has to place the desired initial value of X_i into the array and pass it as an
 8768 argument. By using different arguments, *erand48()*, *rand48()*, and *jrand48()* allow separate
 8769 modules of a large program to generate several *independent* streams of pseudo-random numbers;
 8770 that is, the sequence of numbers in each stream shall *not* depend upon how many times the
 8771 routines are called to generate numbers for the other streams.

8772 The initializer function *srand48()* sets the high-order 32 bits of X_i to the low-order 32 bits
 8773 contained in its argument. The low-order 16 bits of X_i are set to the arbitrary value $330E_{16}$.

8774 The initializer function *seed48()* sets the value of X_i to the 48-bit value specified in the argument
 8775 array. The low-order 16 bits of X_i are set to the low-order 16 bits of *seed16v*[0]. The mid-order 16
 8776 bits of X_i are set to the low-order 16 bits of *seed16v*[1]. The high-order 16 bits of X_i are set to the
 8777 low-order 16 bits of *seed16v*[2]. In addition, the previous value of X_i is copied into a 48-bit
 8778 internal buffer, used only by *seed48()*, and a pointer to this buffer is the value returned by
 8779 *seed48()*. This returned pointer, which can just be ignored if not needed, is useful if a program is
 8780 to be restarted from a given point at some future time—use the pointer to get at and store the
 8781 last X_i value, and then use this value to re-initialize via *seed48()* when the program is restarted.

8782 The initializer function *lcg48()* allows the user to specify the initial X_i , the multiplier value a ,
 8783 and the addend value c . Argument array elements *param*[0-2] specify X_i , *param*[3-5] specify the
 8784 multiplier a , and *param*[6] specifies the 16-bit addend c . After *lcg48()* is called, a subsequent
 8785 call to either *srand48()* or *seed48()* shall restore the standard multiplier and addend values, a and
 8786 c , specified above.

8787 The *drand48()*, *lrnd48()*, and *mrnd48()* functions need not be reentrant. A function that is not
 8788 required to be reentrant is not required to be thread-safe.

8789 RETURN VALUE

8790 As described in the DESCRIPTION above.

8791 ERRORS

8792 No errors are defined.

8793 EXAMPLES

8794 None.

8795 APPLICATION USAGE

8796 None.

8797 RATIONALE

8798 None.

8799 FUTURE DIRECTIONS

8800 None.

8801 SEE ALSO

8802 *rand()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

8803 CHANGE HISTORY

8804 First released in Issue 1. Derived from Issue 1 of the SVID.

8805 Issue 4

8806 The type **long** is replaced by **long** and the type **unsigned short** is replaced by **unsigned short** in
 8807 the SYNOPSIS section.

8808 In the DESCRIPTION, the description of *srand48()* is amended to fix a limitation in Issue 3,
 8809 which indicates that the high-order 32 bits of X_i are set to the {LONG_BIT} bits in the argument.
 8810 Though unintentional, the implication of this statement is that {LONG_BIT} would be 32 on all

- 8811 systems compliant with Issue 3, when in fact Issue 3 imposes no such restriction.
- 8812 The `<stdlib.h>` header is added to the SYNOPSIS section.
- 8813 The argument list for the `lrnd48()` and `mrnd48()` functions now contains **void**.
- 8814 **Issue 5**
- 8815 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.
- 8816 **Issue 6**
- 8817 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

8818 **NAME**

8819 dup, dup2 — duplicate an open file descriptor

8820 **SYNOPSIS**

8821 #include <unistd.h>

8822 int dup(int *fil-des*);8823 int dup2(int *fil-des*, int *fil-des2*);8824 **DESCRIPTION**8825 The *dup()* and *dup2()* functions provide an alternative interface to the service provided by
8826 *fcntl()* using the `F_DUPFD` command. The call:8827 *fid* = dup(*fil-des*);

8828 is equivalent to:

8829 *fid* = fcntl(*fil-des*, `F_DUPFD`, 0);

8830 The call:

8831 *fid* = dup2(*fil-des*, *fil-des2*);

8832 is equivalent to:

8833 close(*fil-des2*);8834 *fid* = fcntl(*fil-des*, `F_DUPFD`, *fil-des2*);

8835 except for the following:

- 8836 • If *fil-des2* is less than 0 or greater than or equal to `{OPEN_MAX}`, *dup2()* shall return `-1` with
8837 *errno* set to `[EBADF]`.
- 8838 • If *fil-des* is a valid file descriptor and is equal to *fil-des2*, *dup2()* shall return *fil-des2* without
8839 closing it.
- 8840 • If *fil-des* is not a valid file descriptor, *dup2()* shall return `-1` and shall not close *fil-des2*.
- 8841 • The value returned shall be equal to the value of *fil-des2* upon successful completion, or `-1`
8842 upon failure.

8843 **RETURN VALUE**8844 Upon successful completion a non-negative integer, namely the file descriptor, shall be returned;
8845 otherwise, `-1` shall be returned and *errno* set to indicate the error.8846 **ERRORS**8847 The *dup()* function shall fail if:8848 `[EBADF]` The *fil-des* argument is not a valid open file descriptor.8849 `[EMFILE]` The number of file descriptors in use by this process would exceed
8850 `{OPEN_MAX}`.8851 The *dup2()* function shall fail if:8852 `[EBADF]` The *fil-des* argument is not a valid open file descriptor or the argument *fil-des2* is
8853 negative or greater than or equal to `{OPEN_MAX}`.8854 `[EINTR]` The *dup2()* function was interrupted by a signal.

8855 **EXAMPLES**8856 **Redirecting Standard Output to a File**

8857 The following example closes standard output for the current processes, re-assigns standard
8858 output to go to the file referenced by *pfid*, and closes the original file descriptor to clean up.

```
8859 #include <unistd.h>
8860 ...
8861 int pfd;
8862 ...
8863 close(1);
8864 dup(pfd);
8865 close(pfd);
8866 ...
```

8867 **Redirecting Error Messages**

8868 The following example redirects messages from *stderr* to *stdout*.

```
8869 #include <unistd.h>
8870 ...
8871 dup2(1, 2);
8872 ...
```

8873 **APPLICATION USAGE**

8874 None.

8875 **RATIONALE**

8876 The *dup()* and *dup2()* functions are redundant. Their services are also provided by the *fcntl()*
8877 function. They have been included in this volume of IEEE Std. 1003.1-200x primarily for
8878 historical reasons, since many existing applications use them.

8879 While the brief code segment shown is very similar in behavior to *dup2()*, a conforming
8880 implementation based on other functions defined in this volume of IEEE Std. 1003.1-200x is
8881 significantly more complex. Least obvious is the possible effect of a signal-catching function that
8882 could be invoked between steps and allocate or deallocate file descriptors. This could be avoided
8883 by blocking signals.

8884 The *dup2()* function is not marked obsolescent because it presents a type-safe version of
8885 functionality provided in a type-unsafe version by *fcntl()*. It is used in the POSIX Ada binding.

8886 The *dup2()* function is not intended for use in critical regions as a synchronization mechanism.

8887 In the description of [EBADF], the case of *fildes* being out of range is covered by the given case of
8888 *fildes* not being valid. The descriptions for *fildes* and *fildes2* are different because the only kind of
8889 invalidity that is relevant for *fildes2* is whether it is out of range; that is, it does not matter
8890 whether *fildes2* refers to an open file when the *dup2()* call is made.

8891 **FUTURE DIRECTIONS**

8892 None.

8893 **SEE ALSO**

8894 *close()*, *fcntl()*, *open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<unistd.h>**

8895 **CHANGE HISTORY**

8896 First released in Issue 1. Derived from Issue 1 of the SVID.

8897 **Issue 4**8898 The <**unistd.h**> header is added to the SYNOPSIS section.8899 [EINTR] is no longer required for *dup()* because *fcntl()* does not return [EINTR] for F_DUPFD.

8900 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

8901 • In the DESCRIPTION, the fourth bullet item describing differences between *dup()* and
8902 *dup2()* is added.8903 • In the ERRORS section, error values returned by *dup()* and *dup2()* are now described
8904 separately.8905 **Issue 6**

8906 The RATIONALE section is added.

8907 **NAME**8908 ecvt, fcvt, gcvt — convert a floating-point number to a string (**LEGACY**)8909 **SYNOPSIS**8910 `XSI` #include <stdlib.h>8911 char *ecvt(double value, int ndigit, int *restrict decpt,
8912 int *restrict sign);8913 char *fcvt(double value, int ndigit, int *restrict decpt,
8914 int *restrict sign);

8915 char *gcvt(double value, int ndigit, char *buf);

8916

8917 **DESCRIPTION**8918 The *ecvt()*, *fcvt()*, and *gcvt()* functions shall convert floating-point numbers to null-terminated strings.8920 The *ecvt()* function shall convert *value* to a null-terminated string of *ndigit* digits (where *ndigit* is reduced to an unspecified limit determined by the precision of a **double**) and return a pointer to the string. The high-order digit is non-zero, unless the value is 0. The low-order digit is rounded. The position of the radix character relative to the beginning of the string is stored in the integer pointed to by *decpt* (negative means to the left of the returned digits). If *value* is zero, it is unspecified whether the integer pointed to by *decpt* would be 0 or 1. The radix character is not included in the returned string. If the sign of the result is negative, the integer pointed to by *sign* is non-zero; otherwise, it is 0.

8928 If the converted value is out of range or is not representable, the contents of the returned string are unspecified.

8930 The *fcvt()* function is identical to *ecvt()* except that *ndigit* specifies the number of digits desired after the radix character. The total number of digits in the result string is restricted to an unspecified limit as determined by the precision of a **double**.8933 The *gcvt()* function shall convert *value* to a null-terminated string (similar to that of the *%g* format of *printf()*) in the array pointed to by *buf* and return *buf*. It produces *ndigit* significant digits (limited to an unspecified value determined by the precision of a **double**) in *%f* if possible, or *%e* (scientific notation) otherwise. A minus sign is included in the returned string if *value* is less than 0. A radix character is included in the returned string if *value* is not a whole number. Trailing zeros are suppressed where *value* is not a whole number. The radix character is determined by the current locale. If *setlocale()* has not been called successfully, the default locale, POSIX, is used. The default locale specifies a period (‘.’) as the radix character. The *LC_NUMERIC* category determines the value of the radix character within the current locale.

8942 These functions need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

8944 **RETURN VALUE**8945 The *ecvt()* and *fcvt()* functions shall return a pointer to a null-terminated string of digits.8946 The *gcvt()* function shall return *buf*.8947 The return values from *ecvt()* and *fcvt()* may point to static data which may be overwritten by subsequent calls to these functions.8949 **ERRORS**

8950 No errors are defined.

8951 **EXAMPLES**

8952 None.

8953 **APPLICATION USAGE**8954 *sprintf()* is preferred over this function.8955 **RATIONALE**

8956 None.

8957 **FUTURE DIRECTIONS**

8958 These functions may be withdrawn in a future version.

8959 **SEE ALSO**8960 *printf()*, *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>8961 **CHANGE HISTORY**

8962 First released in Issue 4, Version 2.

8963 **Issue 5**

8964 Moved from X/OPEN UNIX extension to BASE.

8965 Normative text previously in the APPLICATION USAGE section is moved to the
8966 DESCRIPTION.

8967 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

8968 **Issue 6**

8969 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

8970 This function is marked LEGACY.

8971 The **restrict** keyword is added to the *ecvt()* and *fcvt()* prototypes for alignment with the
8972 ISO/IEC 9899:1999 standard.

8973 **NAME**8974 encrypt — encoding function (**CRYPT**)8975 **SYNOPSIS**8976 XSI `#include <unistd.h>`8977 `void encrypt(char block[64], int edflag);`

8978

8979 **DESCRIPTION**

8980 The `encrypt()` function provides (rather primitive) access to an implementation-defined
8981 encoding algorithm. The key generated by `setkey()` is used to encrypt the string `block` with
8982 `encrypt()`.

8983 The `block` argument to `encrypt()` is an array of length 64 bytes containing only the bytes with
8984 numerical value of 0 and 1. The array is modified in place to a similar array using the key set by
8985 `setkey()`. If `edflag` is 0, the argument is encoded. If `edflag` is 1, the argument may be decoded (see
8986 the APPLICATION USAGE section); if the argument is not decoded, `errno` shall be set to
8987 [ENOSYS].

8988 The `encrypt()` function shall not change the setting of `errno` if successful. An application wishing
8989 to check for error situations should set `errno` to 0 before calling `encrypt()`. If `errno` is non-zero on
8990 return, an error has occurred.

8991 The `encrypt()` function need not be reentrant. A function that is not required to be reentrant is
8992 not required to be thread-safe.

8993 **RETURN VALUE**8994 The `encrypt()` function shall return no value.8995 **ERRORS**8996 The `encrypt()` function shall fail if:

8997 [ENOSYS] The functionality is not supported on this implementation.

8998 **EXAMPLES**

8999 None.

9000 **APPLICATION USAGE**

9001 In some environments, decoding might not be implemented. This is related to U.S. Government
9002 restrictions on encryption and decryption routines: the DES decryption algorithm cannot be
9003 exported outside the U.S. Historical practice has been to ship a different version of the
9004 encryption library without the decryption feature in the routines supplied. Thus the exported
9005 version of `encrypt()` does encoding but not decoding.

9006 **RATIONALE**

9007 None.

9008 **FUTURE DIRECTIONS**

9009 None.

9010 **SEE ALSO**9011 `crypt()`, `setkey()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<unistd.h>`9012 **CHANGE HISTORY**

9013 First released in Issue 1. Derived from Issue 1 of the SVID.

9014 **Issue 4**

9015 The <**unistd.h**> header is added to the SYNOPSIS section.

9016 The DESCRIPTION is amended:

- 9017 • To specify the encoding algorithm as implementation-defined
- 9018 • To change entry to function
- 9019 • To make decoding optional

9020 The APPLICATION USAGE section is expanded to explain the restrictions on the availability of
9021 the DES decryption algorithm.

9022 **Issue 5**

9023 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

9024 **Issue 6**

9025 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

9026 **NAME**

9027 endgrent, getgrent, setgrent — group database entry functions

9028 **SYNOPSIS**

```
9029 xSI      #include <grp.h>
          void endgrent(void);
          struct group *getgrent(void);
          void setgrent(void);
9033
```

9034 **DESCRIPTION**9035 The *endgrent()* function may be called to close the group database when processing is complete.

9036 An implementation that provides extended security controls may impose further
 9037 implementation-defined restrictions on accessing the group database. In particular, the system
 9038 may deny the existence of some or all of the group database entries associated with groups other
 9039 than those groups associated with the caller and may omit users other than the caller from the
 9040 list of members of groups in database entries that are returned.

9041 The *getgrent()* function shall return a pointer to a structure containing the broken-out fields of an
 9042 entry in the group database. When first called, *getgrent()* shall return a pointer to a **group**
 9043 structure containing the first entry in the group database. Thereafter, it shall return a pointer to a
 9044 **group** structure containing the next group structure in the group database, so successive calls
 9045 may be used to search the entire database.

9046 The *setgrent()* function effectively rewinds the group database to allow repeated searches.

9047 These functions need not be reentrant. A function that is not required to be reentrant is not
 9048 required to be thread-safe.

9049 **RETURN VALUE**

9050 When first called, *getgrent()* shall return a pointer to the first group structure in the group
 9051 database. Upon subsequent calls it shall return the next group structure in the group database.
 9052 The *getgrent()* function shall return a null pointer on end-of-file or an error and *errno* may be set
 9053 to indicate the error.

9054 The return value may point to a static area which is overwritten by a subsequent call to
 9055 *getgrgid()*, *getgrnam()*, or *getgrent()*.

9056 **ERRORS**9057 The *getgrent()* function may fail if:

- | | | |
|------|----------|--|
| 9058 | [EINTR] | A signal was caught during the operation. |
| 9059 | [EIO] | An I/O error has occurred. |
| 9060 | [EMFILE] | {OPEN_MAX} file descriptors are currently open in the calling process. |
| 9061 | [ENFILE] | The maximum allowable number of files is currently open in the system. |

9062 **EXAMPLES**

9063 None.

9064 **APPLICATION USAGE**

9065 These functions are provided due to their historical usage. Applications should avoid
9066 dependencies on fields in the group database, whether the database is a single file, or where in
9067 the file system name space the database resides. Applications should use *getgrnam()* and
9068 *getgrgid()* whenever possible because it avoids these dependencies.

9069 **RATIONALE**

9070 None.

9071 **FUTURE DIRECTIONS**

9072 None.

9073 **SEE ALSO**

9074 *getgrgid()*, *getgrnam()*, *getlogin()*, *getpwent()*, the Base Definitions volume of
9075 IEEE Std. 1003.1-200x, <grp.h>

9076 **CHANGE HISTORY**

9077 First released in Issue 4, Version 2.

9078 **Issue 5**

9079 Moved from X/OPEN UNIX extension to BASE.

9080 Normative text previously in the APPLICATION USAGE section is moved to the RETURN
9081 VALUE section.

9082 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

9083 **Issue 6**

9084 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

9085 **NAME**

9086 endhostent, freehostent, gethostent, sethostent — network host database functions

9087 **SYNOPSIS**

```
9088     #include <netdb.h>
9089
9089     void endhostent(void);
9090     void freehostent(struct hostent *ptr);
9091     struct hostent *gethostent(void);
9092     void sethostent(int stayopen);
```

9093 **DESCRIPTION**

9094 These functions enable applications to retrieve information about hosts. This information is
 9095 considered to be stored in a database that can be accessed sequentially or randomly.
 9096 Implementation of this database is unspecified.

9097 **Note:** In many cases it is implemented by the Domain Name System, as documented in
 9098 RFC 1034, RFC 1035, and RFC 1886.

9099 Entries are returned in **hostent** structures. Refer to *gethostbyaddr()* for a definition of the **hostent**
 9100 structure.

9101 The *endhostent()* function shall close the connection to the database, releasing any open file
 9102 descriptor.

9103 The *freehostent()* function shall free the memory occupied by the **hostent** structure pointed to by
 9104 *ptr* and any structures pointed to from that structure, provided that **hostent** was obtained by a
 9105 call to *getipnodebyaddr()* or *getipnodebyname()*. Applications shall not call *freehostent()* except to
 9106 pass it a pointer that was obtained from *getipnodebyaddr()* or *getipnodebyname()*.

9107 The *gethostent()* function shall read the next entry in the database, opening and closing a
 9108 connection to the database as necessary.

9109 The *sethostent()* function shall open a connection to the database and set the next entry for
 9110 retrieval to the first entry in the database. If the *stayopen* argument is non-zero, the connection
 9111 shall not be closed by a call to *gethostent()*, *getipnodebyname()*, *gethostbyname()*, *getipnodebyaddr()*,
 9112 or *gethostbyaddr()*, and the implementation may maintain an open file descriptor.

9113 These functions need not be reentrant. A function that is not required to be reentrant is not
 9114 required to be thread-safe.

9115 **RETURN VALUE**

9116 Upon successful completion, the *gethostent()* function shall return a pointer to a **hostent**
 9117 structure if the requested entry was found, and a null pointer if the end of the database was
 9118 reached or the requested entry was not found.

9119 **ERRORS**

9120 No errors are defined for *endhostent()*, *gethostent()*, and *sethostent()*.

9121 **EXAMPLES**

9122 None.

9123 **APPLICATION USAGE**

9124 The **hostent** structure returned by *getipnodebyaddr()* and *getipnodebyname()*, and any structures
 9125 pointed to from those structures, are dynamically allocated. Applications should call
 9126 *freehostent()* to free the memory used by these structures.

9127 The *gethostent()* function may return pointers to static data, which may be overwritten by
 9128 subsequent calls to any of these functions. Applications shall not call *freehostent()* for this area.

9129 **RATIONALE**

9130 None.

9131 **FUTURE DIRECTIONS**

9132 None.

9133 **SEE ALSO**

9134 *endservent()*, *gethostbyaddr()*, *gethostbyname()*, *getipnodebyaddr()*, *getipnodebyname()*, the Base |
9135 Definitions volume of IEEE Std. 1003.1-200x, <**netdb.h**> |

9136 **CHANGE HISTORY**

9137 First released in Issue 6. Derived from the XNS, Issue 5.2 specification. |

9138 **NAME**

9139 endnetent, getnetbyaddr, getnetbyname, getnetent, setnetent — network database functions

9140 **SYNOPSIS**

9141 #include <netdb.h>

9142 void endnetent(void);

9143 struct netent *getnetbyaddr(uint32_t net, int type);

9144 struct netent *getnetbyname(const char *name);

9145 struct netent *getnetent(void);

9146 void setnetent(int stayopen);

9147 **DESCRIPTION**9148 These functions enable applications to retrieve information about networks. This information is
9149 considered to be stored in a database that can be accessed sequentially or randomly.
9150 Implementation of this database is unspecified.9151 The *endnetent()* function shall close the database, releasing any open file descriptor.9152 The *getnetbyaddr()* function shall search the database from the beginning, and find the first entry
9153 for which the address family specified by *type* matches the *n_addrtype* member and the network
9154 number *net* matches the *n_net* member, opening a connection to the database if necessary. The
9155 *net* argument shall be the network number in host byte order.9156 The *getnetbyname()* function shall search the database from the beginning and find the first entry
9157 for which the network name specified by *name* matches the *n_name* member, opening and
9158 closing a connection to the database as necessary.9159 The *getnetent()* function shall read the next entry of the database, opening a connection to the
9160 database if necessary.9161 The *setnetent()* function shall open and rewind the database. If the *stayopen* argument is non-
9162 zero, the connection to the *net* database shall not be closed after each call to *getnetent()* (either
9163 directly, or indirectly through one of the other *getnet**(*)* functions), and the implementation may
9164 maintain an open file descriptor to the database.9165 The *getnetbyaddr()*, *getnetbyname()*, and *getnetent()*, functions shall each return a pointer to a
9166 **netent** structure, the members of which shall contain the fields of an entry in the network
9167 database.9168 These functions need not be reentrant. A function that is not required to be reentrant is not
9169 required to be thread-safe.9170 **RETURN VALUE**9171 Upon successful completion, *getnetbyaddr()*, *getnetbyname()*, and *getnetent()*, shall return a
9172 pointer to a **netent** structure if the requested entry was found, and a null pointer if the end of the
9173 database was reached or the requested entry was not found. Otherwise, a null pointer shall be
9174 returned.9175 **ERRORS**

9176 No errors are defined.

9177 **EXAMPLES**

9178 None.

9179 **APPLICATION USAGE**9180 The *getnetbyaddr()*, *getnetbyname()*, and *getnetent()*, functions may return pointers to static data,
9181 which may be overwritten by subsequent calls to any of these functions.9182 **RATIONALE**

9183 None.

9184 **FUTURE DIRECTIONS**

9185 None.

9186 **SEE ALSO**9187 The Base Definitions volume of IEEE Std. 1003.1-200x, <**netdb.h**>9188 **CHANGE HISTORY**

9189 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

9190 **NAME**

9191 endprotoent, getprotobyname, getprotobynumber, getprotoent, setprotoent — network protocol
 9192 database functions

9193 **SYNOPSIS**

9194 #include <netdb.h>

```
9195 void endprotoent(void);
9196 struct protoent *getprotobyname(const char *name);
9197 struct protoent *getprotobynumber(int proto);
9198 struct protoent *getprotoent(void);
9199 void setprotoent(int stayopen);
```

9200 **DESCRIPTION**

9201 These functions enable applications to retrieve information about protocols. This information is
 9202 considered to be stored in a database that can be accessed sequentially or randomly.
 9203 Implementation of this database is unspecified.

9204 The *endprotoent()* function shall close the connection to the database, releasing any open file
 9205 descriptor.

9206 The *getprotobyname()* function shall search the database from the beginning and find the first
 9207 entry for which the protocol name specified by *name* matches the *p_name* member, opening a
 9208 connection to the database if necessary.

9209 The *getprotobynumber()* function shall search the database from the beginning and find the first
 9210 entry for which the protocol number specified by *proto* matches the *p_proto* member, opening a
 9211 connection to the database if necessary.

9212 The *getprotoent()* function shall read the next entry of the database, opening and closing a
 9213 connection to the database as necessary.

9214 The *setprotoent()* function shall open a connection to the database, and set the next entry to the
 9215 first entry. If the *stayopen* argument is non-zero, the connection to the network protocol database
 9216 shall not be closed after each call to *getprotoent()* (either directly, or indirectly through one of the
 9217 other *getproto**() functions), and the implementation may maintain an open file descriptor for
 9218 the database.

9219 The *getprotobyname()*, *getprotobynumber()*, and *getprotoent()*, functions shall each return a pointer
 9220 to a **protoent** structure, the members of which shall contain the fields of an entry in the network
 9221 protocol database.

9222 These functions need not be reentrant. A function that is not required to be reentrant is not
 9223 required to be thread-safe.

9224 **RETURN VALUE**

9225 Upon successful completion, *getprotobyname()*, *getprotobynumber()*, and *getprotoent()* return a
 9226 pointer to a **protoent** structure if the requested entry was found, and a null pointer if the end of
 9227 the database was reached or the requested entry was not found. Otherwise, a null pointer is
 9228 returned.

9229 **ERRORS**

9230 No errors are defined.

9231 **EXAMPLES**

9232 None.

9233 **APPLICATION USAGE**9234 The *getprotobyname()*, *getprotobynumber()*, and *getprotoent()* functions may return pointers to
9235 static data, which may be overwritten by subsequent calls to any of these functions.9236 **RATIONALE**

9237 None.

9238 **FUTURE DIRECTIONS**

9239 None.

9240 **SEE ALSO**9241 The Base Definitions volume of IEEE Std. 1003.1-200x, <**netdb.h**>9242 **CHANGE HISTORY**

9243 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

9244 **NAME**

9245 endpwent, getpwent, setpwent — user database functions

9246 **SYNOPSIS**

```
9247 xSI #include <pwd.h>
9248 void endpwent(void);
9249 struct passwd *getpwent(void);
9250 void setpwent(void);
9251
```

9252 **DESCRIPTION**9253 The *endpwent()* function may be called to close the user database when processing is complete.

9254 An implementation that provides extended security controls may impose further
 9255 implementation-defined restrictions on accessing the user database. In particular, the system
 9256 may deny the existence of some or all of the user database entries associated with users other
 9257 than the caller.

9258 The *getpwent()* function shall return a pointer to a structure containing the broken-out fields of
 9259 an entry in the user database. Each entry in the user database contains a **passwd** structure. When
 9260 first called, *getpwent()* shall return a pointer to a **passwd** structure containing the first entry in
 9261 the user database. Thereafter, it shall return a pointer to a **passwd** structure containing the next
 9262 entry in the user database. Successive calls can be used to search the entire user database.

9263 If an end-of-file or an error is encountered on reading, *getpwent()* shall return a null pointer.9264 The *setpwent()* function effectively rewinds the user database to allow repeated searches.

9265 These functions need not be reentrant. A function that is not required to be reentrant is not
 9266 required to be thread-safe.

9267 **RETURN VALUE**9268 The *getpwent()* function shall return a null pointer on end-of-file or error.9269 **ERRORS**9270 The *getpwent()*, *setpwent()*, and *endpwent()* functions may fail if:

9271 [EIO] An I/O error has occurred.

9272 In addition, *getpwent()* and *setpwent()* may fail if:

9273 [EMFILE] {OPEN_MAX} file descriptors are currently open in the calling process.

9274 [ENFILE] The maximum allowable number of files is currently open in the system.

9275 The return value may point to a static area which is overwritten by a subsequent call to
 9276 *getpwuid()*, *getpwnam()*, or *getpwent()*.

9277 **EXAMPLES**9278 **Searching the User Database**

9279 The following example uses the *getpwent()* function to get successive entries in the user
 9280 database, returning a pointer to a **passwd** structure that contains information about each user.
 9281 The call to *endpwent()* closes the user database and cleans up.

```
9282 #include <pwd.h>
9283 ...
9284 struct passwd *p;
9285 ...
```

```
9286     while ((p = getpwent ()) != NULL) {
9287         ...
9288     }
9289     endpwent();
9290     ...
```

9291 APPLICATION USAGE

9292 These functions are provided due to their historical usage. Applications should avoid
9293 dependencies on fields in the password database, whether the database is a single file, or where
9294 in the file system name space the database resides. Applications should use *getpwuid()*
9295 whenever possible because it avoids these dependencies.

9296 RATIONALE

9297 None.

9298 FUTURE DIRECTIONS

9299 None.

9300 SEE ALSO

9301 *endgrent()*, *getlogin()*, *getpwnam()*, *getpwuid()*, the Base Definitions volume of
9302 IEEE Std. 1003.1-200x, <pwd.h>

9303 CHANGE HISTORY

9304 First released in Issue 4, Version 2.

9305 Issue 5

9306 Moved from X/OPEN UNIX extension to BASE.

9307 Normative text previously in the APPLICATION USAGE section is moved to the RETURN
9308 VALUE section.

9309 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

9310 Issue 6

9311 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

9312 **NAME**

9313 endservent, getservbyname, getservbyport, getservent, setservent — network services database
 9314 functions

9315 **SYNOPSIS**

```
9316     #include <netdb.h>

9317     void endservent(void);
9318     struct servent *getservbyname(const char *name, const char *proto);
9319     struct servent *getservbyport(int port, const char *proto);
9320     struct servent *getservent(void);
9321     void setservent(int stayopen);
```

9322 **DESCRIPTION**

9323 These functions enable applications to retrieve information about network services. This
 9324 information is considered to be stored in a database that can be accessed sequentially or
 9325 randomly. Implementation of this database is unspecified.

9326 The *endservent()* function shall close the database, releasing any open file descriptor.

9327 The *getservbyname()* function shall search the database from the beginning and find the first
 9328 entry for which the service name specified by *name* matches the *s_name* member and the protocol
 9329 name specified by *proto* matches the *s_proto* member, opening a connection to the database if
 9330 necessary. If *proto* is a null pointer, any value of the *s_proto* member shall be matched.

9331 The *getservbyport()* function shall search the database from the beginning and find the first entry
 9332 for which the port specified by *port* matches the *s_port* member and the protocol name specified
 9333 by *proto* matches the *s_proto* member, opening a connection to the database if necessary. If *proto*
 9334 is a null pointer, any value of the *s_proto* member shall be matched. The *port* argument shall be in
 9335 network byte order.

9336 The *getservent()* function shall read the next entry of the database, opening and closing a
 9337 connection to the database as necessary.

9338 The *setservent()* function shall open a connection to the database, and set the next entry to the
 9339 first entry. If the *stayopen* argument is non-zero, the *net* database shall not be closed after each
 9340 call to the *getservent()* function (either directly, or indirectly through one of the other *getserv*()*
 9341 functions), and the implementation may maintain an open file descriptor for the database.

9342 The *getservbyname()*, *getservbyport()*, and *getservent()* functions shall each return a pointer to a
 9343 **servent** structure, the members of which shall contain the fields of an entry in the network
 9344 services database.

9345 These functions need not be reentrant. A function that is not required to be reentrant is not
 9346 required to be thread-safe.

9347 **RETURN VALUE**

9348 Upon successful completion, *getservbyname()*, *getservbyport()*, and *getservent()* return a pointer to
 9349 a **servent** structure if the requested entry was found, and a null pointer if the end of the database
 9350 was reached or the requested entry was not found. Otherwise, a null pointer is returned.

9351 **ERRORS**

9352 No errors are defined.

9353 **EXAMPLES**

9354 None.

9355 **APPLICATION USAGE**9356 The *port* argument of *getservbyport()* need not be compatible with the port values of all address
9357 families.9358 The *getservbyname()*, *getservbyport()*, and *getservent()* functions may return pointers to static
9359 data, which may be overwritten by subsequent calls to any of these functions.9360 **RATIONALE**

9361 None.

9362 **FUTURE DIRECTIONS**

9363 None.

9364 **SEE ALSO**9365 *endhostent()*, *endprotoent()*, *htonl()*, *inet_addr()*, the Base Definitions volume of |
9366 IEEE Std. 1003.1-200x, <**netdb.h**> |9367 **CHANGE HISTORY**

9368 First released in Issue 6. Derived from the XNS, Issue 5.2 specification. |

9369 **NAME**

9370 endutxent, getutxent, getutxid, getutxline, pututxline, setutxent — user accounting database
 9371 functions

9372 **SYNOPSIS**

```
9373 xSI #include <utmpx.h>
9374
9374 void endutxent(void);
9375 struct utmpx *getutxent(void);
9376 struct utmpx *getutxid(const struct utmpx *id);
9377 struct utmpx *getutxline(const struct utmpx *line);
9378 struct utmpx *pututxline(const struct utmpx *utmpx);
9379 void setutxent(void);
9380
```

9381 **DESCRIPTION**

9382 These functions provide access to the user accounting database.

9383 The *endutxent()* function closes the user accounting database.

9384 An implementation that provides extended security controls may impose further
 9385 implementation-defined restrictions on accessing the user accounting database. In particular, the
 9386 system may deny the existence of some or all of the user accounting database entries associated
 9387 with users other than the caller.

9388 The *getutxent()* function reads in the next entry from the user accounting database. If the
 9389 database is not already open, it opens it. If it reaches the end of the database, it fails.

9390 The *getutxid()* function searches forward from the current point in the database. If the *ut_type*
 9391 value of the **utmpx** structure pointed to by *id* is *BOOT_TIME*, *OLD_TIME*, or *NEW_TIME*, then
 9392 it stops when it finds an entry with a matching *ut_type* value. If the *ut_type* value is
 9393 *INIT_PROCESS*, *LOGIN_PROCESS*, *USER_PROCESS*, or *DEAD_PROCESS*, then it stops when
 9394 it finds an entry whose type is one of these four and whose *ut_id* member matches the *ut_id*
 9395 member of the **utmpx** structure pointed to by *id*. If the end of the database is reached without a
 9396 match, *getutxid()* fails.

9397 The *getutxid()* or *getutxline()* function may cache data. For this reason, to use *getutxline()* to
 9398 search for multiple occurrences, it is necessary to zero out the static data after each success, or
 9399 *getutxline()* could just return a pointer to the same **utmpx** structure over and over again.

9400 There is one exception to the rule about removing the structure before further reads are done.
 9401 The implicit read done by *pututxline()* (if it finds that it is not already at the correct place in the
 9402 user accounting database) shall not modify the static structure returned by *getutxent()*,
 9403 *getutxid()*, or *getutxline()*, if the application has just modified this structure and passed the
 9404 pointer back to *pututxline()*.

9405 For all entries that match a request, the *ut_type* member indicates the type of the entry. Other
 9406 members of the entry shall contain meaningful data based on the value of the *ut_type* member as
 9407 follows:

9408
9409
9410
9411
9412
9413
9414
9415
9416
9417
9418

ut_type Member	Other Members with Meaningful Data
EMPTY	No others
BOOT_TIME	<i>ut_tv</i>
OLD_TIME	<i>ut_tv</i>
NEW_TIME	<i>ut_tv</i>
USER_PROCESS	<i>ut_id</i> , <i>ut_user</i> (login name of the user), <i>ut_line</i> , <i>ut_pid</i> , <i>ut_tv</i>
INIT_PROCESS	<i>ut_id</i> , <i>ut_pid</i> , <i>ut_tv</i>
LOGIN_PROCESS	<i>ut_id</i> , <i>ut_user</i> (implementation-defined name of the login process), <i>ut_pid</i> , <i>ut_tv</i>
DEAD_PROCESS	<i>ut_id</i> , <i>ut_pid</i> , <i>ut_tv</i>

9419
9420
9421
9422

The *getutxline()* function searches forward from the current point in the database until it finds an entry of the type LOGIN_PROCESS or USER_PROCESS which also has a *ut_line* value matching that in the **utmpx** structure pointed to by *line*. If the end of the database is reached without a match, *getutxline()* fails.

9423
9424
9425
9426

If the process has appropriate privileges, the *pututxline()* function writes out the structure into the user accounting database. It uses *getutxid()* to search for a record that satisfies the request. If this search succeeds, then the entry is replaced. Otherwise, a new entry is made at the end of the user accounting database.

9427
9428

The *setutxent()* function resets the input to the beginning of the database. This should be done before each search for a new entry if it is desired that the entire database be examined.

9429
9430

These functions need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

9431 RETURN VALUE

9432
9433
9434

Upon successful completion, *getutxent()*, *getutxid()*, and *getutxline()* shall return a pointer to a **utmpx** structure containing a copy of the requested entry in the user accounting database. Otherwise, a null pointer shall be returned.

9435
9436

The return value may point to a static area which is overwritten by a subsequent call to *getutxid()* or *getutxline()*.

9437
9438
9439

Upon successful completion, *pututxline()* shall return a pointer to a **utmpx** structure containing a copy of the entry added to the user accounting database. Otherwise, a null pointer shall be returned.

9440

The *endutxent()* and *setutxent()* functions shall return no value.

9441 ERRORS

9442
9443

No errors are defined for the *endutxent()*, *getutxent()*, *getutxid()*, *getutxline()*, and *setutxent()* functions.

9444

The *pututxline()* function may fail if:

9445

[EPERM] The process does not have appropriate privileges.

9446 **EXAMPLES**

9447 None.

9448 **APPLICATION USAGE**9449 The sizes of the arrays in the structure can be found using the *sizeof* operator.9450 **RATIONALE**

9451 None.

9452 **FUTURE DIRECTIONS**

9453 None.

9454 **SEE ALSO**

9455 The Base Definitions volume of IEEE Std. 1003.1-200x, <utmpx.h>

9456 **CHANGE HISTORY**

9457 First released in Issue 4, Version 2.

9458 **Issue 5**

9459 Moved from X/OPEN UNIX extension to BASE.

9460 Normative text previously in the APPLICATION USAGE section is moved to the
9461 DESCRIPTION.

9462 A note indicating that these functions need not be reentrant is added to the DESCRIPTION.

9463 **Issue 6**

9464 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

9465 **NAME**9466 **environ** — array of character pointers to the environment strings9467 **SYNOPSIS**

9468 extern char **environ;

9469 **DESCRIPTION**9470 Refer to the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 8, Environment Variables
9471 and *exec*. |

9472 **NAME**

9473 erand48 — generate uniformly distributed pseudo-random numbers

9474 **SYNOPSIS**

9475 xSI #include <stdlib.h>

9476 double erand48(unsigned short xsubi[3]);

9477

9478 **DESCRIPTION**9479 Refer to *drand48()*.

9480 **NAME**

9481 erf, erfc, erfcf, erfcl, erff, erfl — error and complementary error functions

9482 **SYNOPSIS**

```

9483 xSI #include <math.h>
9484     double erf(double x);
9485     double erfc(double x);
9486     float erfcf(float x);
9487     long double erfcl(long double x);
9488     float erff(float x);
9489     long double erfl(long double x);
9490

```

9491 **DESCRIPTION**9492 The *erf()*, *erff()* and *erfl()* functions shall compute the error function of *x*, defined as:

$$9493 \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

9494 The *erfc()*, *erfcf()*, and *erfcl()* functions shall compute $1.0 - \text{erf}(x)$.9495 An application wishing to check for error situations should set *errno* to 0 before calling *erf()*. If
9496 *errno* is non-zero on return, or the return value is NaN, an error has occurred.9497 **RETURN VALUE**9498 Upon successful completion, these functions shall return the value of the error function and
9499 complementary error function, respectively.9500 If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].9501 If the correct value would cause underflow, 0 shall be returned and *errno* may be set to
9502 [ERANGE].9503 **ERRORS**9504 The *erfc()*, *erfcf()*, and *erfcl()* functions shall fail if:9505 [ERANGE] The value of *x* is too large.9506 The *erf()* and *erfc()* functions may fail if:9507 [EDOM] The value of *x* is NaN.

9508 [ERANGE] The result underflows

9509 No other errors shall occur.

9510 **EXAMPLES**

9511 None.

9512 **APPLICATION USAGE**9513 The *erfc()* function is provided because of the extreme loss of relative accuracy if *erf(x)* is called
9514 for large *x* and the result subtracted from 1.0.9515 **RATIONALE**

9516 None.

9517 **FUTURE DIRECTIONS**

9518 None.

9519 **SEE ALSO**

9520 *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

9521 **CHANGE HISTORY**

9522 First released in Issue 1. Derived from Issue 1 of the SVID.

9523 **Issue 4**

9524 References to *matherr()* are removed.

9525 The RETURN VALUE and ERRORS sections are substantially rewritten to rationalize error
9526 handling in the mathematics functions.

9527 **Issue 5**

9528 The DESCRIPTION is updated to indicate how an application should check for an error. This
9529 text was previously published in the APPLICATION USAGE section.

9530 **Issue 6**

9531 The *erfcf()*, *erfcl()*, *erff()*, and *erfl()* functions are added for alignment with the
9532 ISO/IEC 9899:1999 standard.

9533 **NAME**

9534 errno — error return value

9535 **SYNOPSIS**

9536 #include <errno.h>

9537 **DESCRIPTION**9538 *errno* is used by many functions to return error values.

9539 Many functions provide an error number in *errno* which has type **int** and is defined in <**errno.h**>. The value of *errno* shall be defined only after a call to a function for which it is explicitly stated to be set and until it is changed by the next function call. The value of *errno* should only be examined when it is indicated to be valid by a function's return value. Programs should obtain the definition of *errno* by the inclusion of <**errno.h**>. No function in this volume of IEEE Std. 1003.1-200x shall set *errno* to 0.

9545 It is unspecified whether *errno* is a macro or an identifier declared with external linkage. If a macro definition is suppressed in order to access an actual object, or a program defines an identifier with the name *errno*, the behavior is undefined.

9548 The symbolic values stored in *errno* are documented in the ERRORS sections on all relevant pages.

9550 **RETURN VALUE**

9551 None.

9552 **ERRORS**

9553 None.

9554 **EXAMPLES**

9555 None.

9556 **APPLICATION USAGE**

9557 Previously both POSIX and X/Open documents were more restrictive than the ISO C standard in that they required *errno* to be defined as an external variable, whereas the ISO C standard required only that *errno* be defined as a modifiable **lvalue** with type **int**.

9560 A program that uses *errno* for error checking should set it to 0 before a function call, then inspect it before a subsequent function call.

9562 **RATIONALE**

9563 None.

9564 **FUTURE DIRECTIONS**

9565 None.

9566 **SEE ALSO**9567 Section 2.3, the Base Definitions volume of IEEE Std. 1003.1-200x, <**errno.h**>9568 **CHANGE HISTORY**

9569 First released in Issue 1. Derived from Issue 1 of the SVID.

9570 **Issue 4**

9571 The FUTURE DIRECTIONS section is deleted.

9572 The following changes are incorporated for alignment with the ISO C standard:

- 9573 • The DESCRIPTION now guarantees that *errno* is set to 0 at program start-up, and that it is never reset to 0 by any XSI function.

9574

- 9575 • The APPLICATION USAGE section is added. This issue is aligned with the ISO C standard,
9576 which permits *errno* to be a macro.

9577 **Issue 5**

9578 The following sentence is deleted from the DESCRIPTION: “The value of *errno* is 0 at program
9579 start-up, but is never set to 0 by any XSI function”. The DESCRIPTION also no longer states that
9580 conforming implementations may support the declaration:

9581 extern int errno;

9582 **Issue 6**

9583 Obsolescent text regarding defining *errno* as:

9584 extern int errno

9585 is removed.

9586 Text regarding no function setting *errno* to zero to indicate an error is changed to no function
9587 shall set *errno* to zero. This is for alignment with the ISO/IEC 9899:1999 standard.

9588 NAME

9589 environ, execl, execv, execl, execve, execlp, execvp — execute a file

9590 SYNOPSIS

```

9591     #include <unistd.h>

9592     extern char **environ;
9593     int execl(const char *path, const char *arg0, ... /*, (char *)0 */);
9594     int execv(const char *path, char *const argv[]);
9595     int execl(const char *path, const char *arg0, ... /*,
9596             (char *)0, char *const envp[] */);
9597     int execve(const char *path, char *const argv[], char *const envp[]);
9598     int execlp(const char *file, const char *arg0, ... /*, (char *)0 */);
9599     int execvp(const char *file, char *const argv[]);

```

9600 DESCRIPTION

9601 The *exec* functions shall replace the current process image with a new process image. The new
 9602 image is constructed from a regular, executable file called the *new process image file*. There shall
 9603 be no return from a successful *exec*, because the calling process image is overlaid by the new
 9604 process image.

9605 When a C-language program is executed as a result of this call, it shall be entered as a C-
 9606 language function call as follows:

```
9607     int main (int argc, char *argv[]);
```

9608 where *argc* is the argument count and *argv* is an array of character pointers to the arguments
 9609 themselves. In addition, the following variable:

```
9610     extern char **environ;
```

9611 is initialized as a pointer to an array of character pointers to the environment strings. The *argv*
 9612 and *environ* arrays are each terminated by a null pointer. The null pointer terminating the *argv*
 9613 array is not counted in *argc*.

9614 THR Conforming multi-threaded applications shall not use the *environ* variable to access or modify
 9615 any environment variable while any other thread is concurrently modifying any environment
 9616 variable. A call to any function dependent on any environment variable shall be considered a use
 9617 of the *environ* variable to access that environment variable.

9618 The arguments specified by a program with one of the *exec* functions shall be passed on to the
 9619 new process image in the corresponding *main()* arguments.

9620 The argument *path* points to a path name that identifies the new process image file.

9621 The argument *file* is used to construct a path name that identifies the new process image file. If
 9622 the *file* argument contains a slash character, the *file* argument shall be used as the path name for
 9623 this file. Otherwise, the path prefix for this file is obtained by a search of the directories passed
 9624 as the environment variable *PATH* (see the Base Definitions volume of IEEE Std. 1003.1-200x,
 9625 Chapter 8, Environment Variables). If this environment variable is not present, the results of the
 9626 search are implementation-defined.

9627 If the process image file is not a valid executable object, *execlp()* and *execvp()* shall use the
 9628 contents of that file as standard input to a command interpreter conforming to *system()*. In this
 9629 case, the command interpreter becomes the new process image.

9630 The arguments represented by *arg0,...* are pointers to null-terminated character strings. These
 9631 strings constitute the argument list available to the new process image. The list is terminated by
 9632 a null pointer. The argument *arg0* should point to a file name that is associated with the process

- 9633 being started by one of the *exec* functions.
- 9634 The argument *argv* is an array of character pointers to null-terminated strings. The application
9635 shall ensure that the last member of this array is a null pointer. These strings constitute the
9636 argument list available to the new process image. The value in *argv*[0] should point to a file
9637 name that is associated with the process being started by one of the *exec* functions.
- 9638 The argument *envp* is an array of character pointers to null-terminated strings. These strings
9639 constitute the environment for the new process image. The *envp* array is terminated by a null
9640 pointer.
- 9641 For those forms not containing an *envp* pointer (*execl*(), *execv*(), *execlp*(), and *execvp*()), the
9642 environment for the new process image is taken from the external variable *environ* in the calling
9643 process.
- 9644 The number of bytes available for the new process' combined argument and environment lists is
9645 {ARG_MAX}. It is implementation-defined whether null terminators, pointers, and/or any
9646 alignment bytes are included in this total.
- 9647 File descriptors open in the calling process image remain open in the new process image, except
9648 for those whose close-on-exec flag FD_CLOEXEC is set. For those file descriptors that remain
9649 open, all attributes of the open file description remain unchanged. For any file descriptor that is
9650 closed for this reason, file locks are removed as a result of the close as described in *close*(). Locks
9651 that are not removed by closing of file descriptors remain unchanged.
- 9652 Directory streams open in the calling process image shall be closed in the new process image.
- 9653 XSI The state of conversion descriptors and message catalog descriptors in the new process image is
9654 undefined. For the new process, the equivalent of:
- ```
9655 setlocale(LC_ALL, "C")
```
- 9656 is executed at start-up.
- 9657 Signals set to the default action (SIG\_DFL) in the calling process image shall be set to the default  
9658 action in the new process image. Signals set to be ignored (SIG\_IGN) by the calling process  
9659 image shall be set to be ignored by the new process image. Signals set to be caught by the calling  
9660 XSI process image shall be set to the default action in the new process image (see <signal.h>). After  
9661 a successful call to any of the *exec* functions, alternate signal stacks are not preserved and the  
9662 SA\_ONSTACK flag shall be cleared for all signals.
- 9663 After a successful call to any of the *exec* functions, any functions previously registered by *atexit*()  
9664 are no longer registered.
- 9665 XSI If the ST\_NOSUID bit is set for the file system containing the new process image file, then the  
9666 effective user ID, effective group ID, saved set-user-ID, and saved set-group-ID are unchanged  
9667 in the new process image. Otherwise, if the set-user-ID mode bit of the new process image file is  
9668 set, the effective user ID of the new process image is set to the user ID of the new process image  
9669 file. Similarly, if the set-group-ID mode bit of the new process image file is set, the effective  
9670 group ID of the new process image is set to the group ID of the new process image file. The real  
9671 user ID, real group ID, and supplementary group IDs of the new process image remain the same  
9672 as those of the calling process image. The effective user ID and effective group ID of the new  
9673 process image shall be saved (as the saved set-user-ID and the saved set-group-ID) for use by  
9674 *setuid*().
- 9675 XSI Any shared memory segments attached to the calling process image shall not be attached to the  
9676 new process image.

|      |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 9677 | SEM    | Any named semaphores open in the calling process shall be closed as if by appropriate calls to <i>sem_close()</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 9678 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9679 | TYM    | Any blocks of typed memory that were mapped in the calling process are unmapped, as if <i>munmap()</i> was implicitly called to unmap them.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| 9680 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9681 | ML     | Memory locks established by the calling process via calls to <i>mlockall()</i> or <i>mlock()</i> shall be removed. If locked pages in the address space of the calling process are also mapped into the address spaces of other processes and are locked by those processes, the locks established by the other processes shall be unaffected by the call by this process to the <i>exec</i> function. If the <i>exec</i> function fails, the effect on memory locks is unspecified.                                                                                                                                                                                       |
| 9682 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9683 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9684 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9685 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9686 | MF SHM | Memory mappings created in the process are unmapped before the address space is rebuilt for the new process image.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 9687 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9688 | PS     | For the <i>SCHED_FIFO</i> and <i>SCHED_RR</i> scheduling policies, the policy and priority settings shall not be changed by a call to an <i>exec</i> function. For other scheduling policies, the policy and priority settings on <i>exec</i> are implementation-defined.                                                                                                                                                                                                                                                                                                                                                                                                  |
| 9689 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9690 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9691 | TMR    | Per-process timers created by the calling process shall be deleted before replacing the current process image with the new process image.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| 9692 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9693 | MSG    | All open message queue descriptors in the calling process shall be closed, as described in <i>mq_close()</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 9694 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9695 | AIO    | Any outstanding asynchronous I/O operations may be canceled. Those asynchronous I/O operations that are not canceled shall complete as if the <i>exec</i> function had not yet occurred, but any associated signal notifications shall be suppressed. It is unspecified whether the <i>exec</i> function itself blocks awaiting such I/O completion. In no event, however, shall the new process image created by the <i>exec</i> function be affected by the presence of outstanding asynchronous I/O operations at the time the <i>exec</i> function is called. Whether any I/O is canceled, and which I/O may be canceled upon <i>exec</i> , is implementation-defined. |
| 9696 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9697 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9698 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9699 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9700 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9701 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9702 | CPT    | The new process image shall inherit the CPU-time clock of the calling process image. This inheritance means that the process CPU-time clock of the process being <i>execed</i> shall not be reinitialized or altered as a result of the <i>exec</i> function other than to reflect the time spent by the process executing the <i>exec</i> function itself.                                                                                                                                                                                                                                                                                                                |
| 9703 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9704 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9705 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9706 | TCT    | The initial value of the CPU-time clock of the initial thread of the new process image shall be set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 9707 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9708 | TRC    | If the calling process is being traced, the new process image continues to be traced into the same trace stream as the original process image, but the new process image shall not inherit the mapping of trace event names to trace event type identifiers that was defined by calls to the <i>posix_trace_eventid_open()</i> or the <i>posix_trace_trid_eventid_open()</i> functions in the calling process image.                                                                                                                                                                                                                                                       |
| 9709 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9710 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9711 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9712 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9713 |        | If the calling process is a trace controller process, any trace streams that were created by the calling process shall be shut down as described in the <i>posix_trace_shutdown()</i> function.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9714 |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9715 |        | The new process also inherits at least the following attributes from the calling process image:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 9716 | XSI    | <ul style="list-style-type: none"> <li>• Nice value (see <i>nice()</i>)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 9717 | XSI    | <ul style="list-style-type: none"> <li>• <i>semadj</i> values (see <i>semop()</i>)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 9718 |        | <ul style="list-style-type: none"> <li>• Process ID</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 9719 |        | <ul style="list-style-type: none"> <li>• Parent process ID</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

- 9720           • Process group ID
  - 9721           • Session membership
  - 9722           • Real user ID
  - 9723           • Real group ID
  - 9724           • Supplementary group IDs
  - 9725           • Time left until an alarm clock signal (see *alarm()*)
  - 9726           • Current working directory
  - 9727           • Root directory
  - 9728           • File mode creation mask (see *umask()*)
  - 9729 XSI       • File size limit (see *ulimit()*)
  - 9730           • Process signal mask (see *sigprocmask()*)
  - 9731           • Pending signal (see *sigpending()*)
  - 9732           • *tms\_utime*, *tms\_stime*, *tms\_cutime*, and *tms\_cstime* (see *times()*)
  - 9733 XSI       • Resource limits
  - 9734 XSI       • Controlling terminal
  - 9735 XSI       • Interval timers
- 9736           All other process attributes defined in this volume of IEEE Std. 1003.1-200x shall be the same in  
 9737           the new and old process images. The inheritance of process attributes not defined by this  
 9738           volume of IEEE Std. 1003.1-200x is implementation-defined.
- 9739           A call to any *exec* function from a process with more than one thread results in all threads being  
 9740           terminated and the new executable image being loaded and executed. No destructor functions  
 9741           shall be called.
- 9742           Upon successful completion, the *exec* functions shall mark for update the *st\_atime* field of the file.  
 9743           If an *exec* function failed but was able to locate the *process image file*, whether the *st\_atime* field is  
 9744           marked for update is unspecified. Should the *exec* function succeed, the process image file shall  
 9745           be considered to have been opened with *open()*. The corresponding *close()* shall be considered  
 9746           to occur at a time after this open, but before process termination or successful completion of a  
 9747           subsequent call to one of the *exec* functions. The *argv[]* and *envp[]* arrays of pointers and the  
 9748           strings to which those arrays point shall not be modified by a call to one of the *exec* functions,  
 9749           except as a consequence of replacing the process image.
- 9750 XSI       The saved resource limits in the new process image are set to be a copy of the process'  
 9751           corresponding hard and soft limits.
- 9752   **RETURN VALUE**
- 9753           If one of the *exec* functions returns to the calling process image, an error has occurred; the return  
 9754           value shall be  $-1$ , and *errno* shall be set to indicate the error.
- 9755   **ERRORS**
- 9756           The *exec* functions shall fail if:
- 9757           [E2BIG]           The number of bytes used by the new process image's argument list and  
 9758                            environment list is greater than the system-imposed limit of {ARG\_MAX}  
 9759                            bytes.

|      |                |                                                                                                              |
|------|----------------|--------------------------------------------------------------------------------------------------------------|
| 9760 | [EACCES]       | Search permission is denied for a directory listed in the new process image                                  |
| 9761 |                | file's path prefix, or the new process image file denies execution permission,                               |
| 9762 |                | or the new process image file is not a regular file and the implementation does                              |
| 9763 |                | not support execution of files of its type.                                                                  |
| 9764 | [EINVAL]       | The new process image file has the appropriate permission and has a                                          |
| 9765 |                | recognized executable binary format, but the system does not support                                         |
| 9766 |                | execution of a file with this format.                                                                        |
| 9767 | [ELOOP]        | A loop exists in symbolic links encountered during resolution of the <i>path</i> or <i>file</i>              |
| 9768 |                | argument.                                                                                                    |
| 9769 | [ENAMETOOLONG] |                                                                                                              |
| 9770 |                | The length of the <i>path</i> or <i>file</i> arguments exceeds {PATH_MAX} or a path name                     |
| 9771 |                | component is longer than {NAME_MAX}.                                                                         |
| 9772 | [ENOENT]       | A component of <i>path</i> or <i>file</i> does not name an existing file or <i>path</i> or <i>file</i> is an |
| 9773 |                | empty string.                                                                                                |
| 9774 | [ENOTDIR]      | A component of the new process image file's path prefix is not a directory.                                  |
| 9775 |                | The <i>exec</i> functions, except for <i>execlp()</i> and <i>execvp()</i> , shall fail if:                   |
| 9776 | [ENOEXEC]      | The new process image file has the appropriate access permission but has an                                  |
| 9777 |                | unrecognized format.                                                                                         |
| 9778 |                | The <i>exec</i> functions may fail if:                                                                       |
| 9779 | [ELOOP]        | More than {SYMLOOP_MAX} symbolic links were encountered during                                               |
| 9780 |                | resolution of the <i>path</i> or <i>file</i> argument.                                                       |
| 9781 | [ENAMETOOLONG] |                                                                                                              |
| 9782 |                | As a result of encountering a symbolic link in resolution of the <i>path</i> argument,                       |
| 9783 |                | the length of the substituted path name string exceeded {PATH_MAX}.                                          |
| 9784 | [ENOMEM]       | The new process image requires more memory than is allowed by the                                            |
| 9785 |                | hardware or system-imposed memory management constraints.                                                    |
| 9786 | [ETXTBSY]      | The new process image file is a pure procedure (shared text) file that is                                    |
| 9787 |                | currently open for writing by some process.                                                                  |

## 9788 EXAMPLES

### 9789 Using `execl()`

9790 The following example executes the `ls` command, specifying the path name of the executable  
 9791 (`/bin/ls`) and using arguments supplied directly to the command to produce single-column  
 9792 output.

```
9793 #include <unistd.h>
9794 int ret;
9795 ...
9796 ret = execl ("/bin/ls", "ls", "-l", NULL);
```

**9797 Using execl()**

9798 The following example is similar to **Using execl()** (on page 794). In addition, it specifies the  
9799 environment for the new process image using the *env* argument.

```
9800 #include <unistd.h>
9801 int ret;
9802 char *env[] = { "HOME=/usr/home", "LOGNAME=home", NULL };
9803 ...
9804 ret = execl ("/bin/ls", "ls", "-l", NULL, env);
```

**9805 Using execlp()**

9806 The following example searches for the location of the *ls* command among the directories  
9807 specified by the *PATH* environment variable.

```
9808 #include <unistd.h>
9809 int ret;
9810 ...
9811 ret = execlp ("ls", "ls", "-l", NULL);
```

**9812 Using execv()**

9813 The following example passes arguments to the *ls* command in the *cmd* array.

```
9814 #include <unistd.h>
9815 int ret;
9816 char *cmd[] = { "ls", "-l", NULL };
9817 ...
9818 ret = execv ("/bin/ls", cmd);
```

**9819 Using execve()**

9820 The following example passes arguments to the *ls* command in the *cmd* array, and specifies the  
9821 environment for the new process image using the *env* argument.

```
9822 #include <unistd.h>
9823 int ret;
9824 char *cmd[] = { "ls", "-l", NULL };
9825 char *env[] = { "HOME=/usr/home", "LOGNAME=home", NULL };
9826 ...
9827 ret = execve ("/bin/ls", cmd, env);
```

**9828 Using execvp()**

9829 The following example searches for the location of the *ls* command among the directories  
9830 specified by the *PATH* environment variable, and passes arguments to the *ls* command in the  
9831 *cmd* array.

```
9832 #include <unistd.h>
9833 int ret;
9834 char *cmd[] = { "ls", "-l", NULL };
9835 ...
```

9836           ret = execvp ("ls", cmd);

#### 9837 APPLICATION USAGE

9838           As the state of conversion descriptors and message catalog descriptors in the new process image  
9839           is undefined, portable applications should not rely on their use and should close them prior to  
9840           calling one of the *exec* functions.

9841           Applications that require other than the default POSIX locale should call *setlocale()* with the  
9842           appropriate parameters to establish the locale of the new process.

9843           The *environ* array should not be accessed directly by the application.

#### 9844 RATIONALE

9845           Early proposals required that the value of *argc* passed to *main()* be “one or greater”. This was  
9846           driven by the same requirement in drafts of the ISO C standard. In fact, historical  
9847           implementations have passed a value of zero when no arguments are supplied to the caller of  
9848           the *exec* functions. This requirement was removed from the ISO C standard and subsequently  
9849           removed from this volume of IEEE Std. 1003.1-200x as well. The wording, in particular the use of  
9850           the word *should*, requires a Strictly Conforming POSIX Application to pass at least one argument  
9851           to the *exec* function, thus guaranteeing that *argc* be one or greater when invoked by such an  
9852           application. In fact, this is good practice, since many existing applications reference *argv[0]*  
9853           without first checking the value of *argc*.

9854           The requirement on a Strictly Conforming POSIX Application also states that the value passed  
9855           as the first argument be a file name associated with the process being started. Although some  
9856           existing applications pass a path name rather than a file name in some circumstances, a file  
9857           name is more generally useful, since the common usage of *argv[0]* is in printing diagnostics. In  
9858           some cases the file name passed is not the actual file name of the file; for example, many  
9859           implementations of the *login* utility use a convention of prefixing a hyphen (‘-’) to the actual  
9860           file name, which indicates to the command interpreter being invoked that it is a “login shell”.

9861           Some systems can *exec* shell scripts.

9862           One common historical implementation is that the *execl()*, *execv()*, *execle()*, and *execve()*  
9863           functions return an [ENOEXEC] error for any file not recognizable as executable, including a  
9864           shell script. When the *execlp()* and *execvp()* functions encounter such a file, they assume the file  
9865           to be a shell script and invoke a known command interpreter to interpret such files. These  
9866           implementations of *execvp()* and *execlp()* only give the [ENOEXEC] error in the rare case of a  
9867           problem with the command interpreter’s executable file. Because of these implementations, the  
9868           [ENOEXEC] error is not mentioned for *execlp()* or *execvp()*, although implementations can still  
9869           give it.

9870           Another way that some historical implementations handle shell scripts is by recognizing the first  
9871           two bytes of the file as the character string “#!” and using the remainder of the first line of the  
9872           file as the name of the command interpreter to execute.

9873           Some implementations provide a third argument to *main()* called *envp*. This is defined as a  
9874           pointer to the environment. The ISO C standard specifies invoking *main()* with two arguments,  
9875           so implementations must support applications written this way. Since this volume of  
9876           IEEE Std. 1003.1-200x defines the global variable *environ*, which is also provided by historical  
9877           implementations and can be used anywhere that *envp* could be used, there is no functional need  
9878           for the *envp* argument. Applications should use the *getenv()* function rather than accessing the  
9879           environment directly via either *envp* or *environ*. Implementations are required to support the  
9880           two-argument calling sequence, but this does not prohibit an implementation from supporting  
9881           *envp* as an optional third argument.

9882 This volume of IEEE Std. 1003.1-200x specifies that signals set to SIG\_IGN remain set to  
 9883 SIG\_IGN, and that the process signal mask be unchanged across an *exec*. This is consistent with  
 9884 historical implementations, and it permits some useful functionality, such as the *nohup*  
 9885 command. However, it should be noted that many existing applications wrongly assume that  
 9886 they start with certain signals set to the default action and/or unblocked. In particular,  
 9887 applications written with a simpler signal model that does not include blocking of signals, such  
 9888 as the one in the ISO C standard, may not behave properly if invoked with some signals blocked.  
 9889 Therefore, it is best not to block or ignore signals across *execs* without explicit reason to do so,  
 9890 and especially not to block signals across *execs* of arbitrary (not closely co-operating) programs.

9891 The *exec* functions always save the value of the effective user ID and effective group ID of the  
 9892 process at the completion of the *exec*, whether or not the set-user-ID or the set-group-ID bit of  
 9893 the process image file is set.

9894 The statement about *argv[]* and *envp[]* being constants is included to make explicit to future  
 9895 writers of language bindings that these objects are completely constant. Due to a limitation of  
 9896 the ISO C standard, it is not possible to state that idea in standard C. Specifying two levels of  
 9897 *const-qualification* for the *argv[]* and *envp[]* parameters for the *exec* functions may seem to be the  
 9898 natural choice, given that these functions do not modify either the array of pointers or the  
 9899 characters to which the function points, but this would disallow existing correct code. Instead,  
 9900 only the array of pointers is noted as constant. The table of assignment compatibility for *dst=src*,  
 9901 derived from the ISO C standard summarizes the compatibility:

| 9902 | <i>dst:</i>         | char *[] | const char *[] | char *const[] | const char *const[] |
|------|---------------------|----------|----------------|---------------|---------------------|
| 9903 | <i>src:</i>         |          |                |               |                     |
| 9904 | char *[]            | VALID    | —              | VALID         | —                   |
| 9905 | const char *[]      | —        | VALID          | —             | VALID               |
| 9906 | char * const []     | —        | —              | VALID         | —                   |
| 9907 | const char *const[] | —        | —              | —             | VALID               |

9908 Since all existing code has a source type matching the first row, the column that gives the most  
 9909 valid combinations is the third column. The only other possibility is the fourth column, but  
 9910 using it would require a cast on the *argv* or *envp* arguments. It is unfortunate that the fourth  
 9911 column cannot be used, because the declaration a non-expert would naturally use would be that  
 9912 in the second row.

9913 The ISO C standard and this volume of IEEE Std. 1003.1-200x do not conflict on the use of  
 9914 *environ*, but some historical implementations of *environ* may cause a conflict. As long as *environ*  
 9915 is treated in the same way as an entry point (for example, *fork()*), it conforms to both standards.  
 9916 A library can contain *fork()*, but if there is a user-provided *fork()*, that *fork()* is given precedence  
 9917 and no problem ensues. The situation is similar for *environ*: the definition in this volume of  
 9918 IEEE Std. 1003.1-200x is to be used if there is no user-provided *environ* to take precedence. At  
 9919 least three implementations are known to exist that solve this problem.

9920 [E2BIG] The limit {ARG\_MAX} applies not just to the size of the argument list, but to  
 9921 the sum of that and the size of the environment list.

9922 [EFAULT] Some historical systems return [EFAULT] rather than [ENOEXEC] when the  
 9923 new process image file is corrupted. They are non-conforming.

9924 [EINVAL] This error condition was added to IEEE Std. 1003.1-200x to allow an  
 9925 implementation to detect executable files generated for different architectures,  
 9926 and indicate this situation to the application. Historical implementations of  
 9927 shells, *execvp()*, and *execlp()* that encounter an [ENOEXEC] error will execute  
 9928 a shell on the assumption that the file is a shell script. This will not produce  
 9929 the desired effect when the file is a valid executable for a different

9930 architecture. An implementation may now choose to avoid this problem by  
 9931 returning [EINVAL] when a valid executable for a different architecture is  
 9932 encountered. Some historical implementations return [EINVAL] to indicate  
 9933 that the *path* argument contains a character with the high order bit set. The  
 9934 standard developers chose to deviate from historical practice for the following  
 9935 reasons:

- 9936 1. The new utilization of [EINVAL] will provide some measure of utility to  
 9937 the user community.
- 9938 2. Historical use of [EINVAL] is not acceptable in an internationalized  
 9939 operating environment.

9940 [ENAMETOOLONG]

9941 Since the file path name may be constructed by taking elements in the *PATH*  
 9942 variable and putting them together with the file name, the  
 9943 [ENAMETOOLONG] error condition could also be reached this way.

9944 [ETXTBSY]

9945 System V returns this error when the executable file is currently open for  
 9946 writing by some process. This volume of IEEE Std. 1003.1-200x neither  
 requires nor prohibits this behavior.

9947 Other systems (such as System V) may return [EINTR] from *exec*. This is not addressed by this  
 9948 volume of IEEE Std. 1003.1-200x, but implementations may have a window between the call to  
 9949 *exec* and the time that a signal could cause one of the *exec* calls to return with [EINTR].

#### 9950 FUTURE DIRECTIONS

9951 None.

#### 9952 SEE ALSO

9953 *alarm()*, *atexit()*, *chmod()*, *close()*, *exit()*, *fcntl()*, *fork()*, *fstatvfs()*, *getenv()*, *getitimer()*, *getrlimit()*,  
 9954 *mmap()*, *nice()*, <REFERENCE UNDEFINED>(posix\_trace\_eventid), *posix\_trace\_shutdown()*,  
 9955 *posix\_trace\_trid\_eventid\_open()*, *putenv()*, *semop()*, *setlocale()*, *shmat()*, *sigaction()*, *sigaltstack()*,  
 9956 *sigpending()*, *sigprocmask()*, *system()*, *times()*, *ulimit()*, *umask()*, the Base Definitions volume of  
 9957 IEEE Std. 1003.1-200x, <**unistd.h**>, the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter  
 9958 11, General Terminal Interface

#### 9959 CHANGE HISTORY

9960 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 9961 Issue 4

9962 The <**unistd.h**> header is added to the SYNOPSIS section.

9963 The **const** keyword is added to identifiers of constant type (for example, *path*, *file*).

9964 In the DESCRIPTION:

- 9965 • An indication of the disposition of conversion descriptors after a call to one of the *exec*  
 9966 functions is added.
- 9967 • A statement about the interaction between *exec* and *atexit()* is added.
- 9968 • *usually* in the descriptions of argument pointers is removed.
- 9969 • *owner ID* is changed to *user ID*.
- 9970 • Shared memory is no longer optional.
- 9971 • The penultimate paragraph is changed to correct an error in Issue 3. It now refers to “All  
 9972 other process attributes ...” instead of “All the process attributes ...”

- 9973 A note about the initialization of locales is added to the APPLICATION USAGE section.
- 9974 The following change is incorporated for alignment with the ISO POSIX-1 standard:
- 9975 • In the ERRORS section, the description of the [ENOEXEC] error is changed to indicate that
  - 9976 this error does not apply to *execlp()* and *execvp()*, and the [ENOMEM] error is added.
- 9977 The following change is incorporated for alignment with the FIPS requirements:
- 9978 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path
  - 9979 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as
  - 9980 an extension.
- 9981 **Issue 4, Version 2**
- 9982 The following changes are incorporated for X/OPEN UNIX conformance:
- 9983 • The DESCRIPTION is changed:
    - 9984 — To indicate the disposition of alternate signal stacks, the SA\_ONSTACK flag, and
    - 9985 mappings established through *mmap()* after a successful call to one of the *exec* functions.
    - 9986 — The effects of ST\_NOSUID being set for a file system are defined.
    - 9987 — A statement is added that mappings established through *mmap()* are not preserved across
    - 9988 an *exec*.
    - 9989 — The list of inherited process attributes is extended to include resource limits, the
    - 9990 controlling terminal, and interval timers.
  - 9991 • In the ERRORS section:
    - 9992 — The condition whereby [ELOOP] is returned if too many symbolic links are encountered
    - 9993 during path name resolution is defined as mandatory.
    - 9994 — A second [ENAMETOOLONG] condition is defined that may report excessive length of
    - 9995 an intermediate result of path name resolution of a symbolic link.
- 9996 **Issue 5**
- 9997 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX
- 9998 Threads Extension.
- 9999 Large File Summit extensions are added.
- 10000 **Issue 6**
- 10001 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:
- 10002 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.
  - 10003 This is since behavior may vary from one file system to another.
- 10004 The following new requirements on POSIX implementations derive from alignment with the
- 10005 Single UNIX Specification:
- 10006 • In the DESCRIPTION, behavior is defined for when the process image file is not a valid
  - 10007 executable.
  - 10008 • In this issue, `_POSIX_SAVED_IDS` is mandated, thus the effective user ID and effective group
  - 10009 ID of the new process image shall be saved (as the saved set-user-ID and the saved set-
  - 10010 group-ID) for use by the *setuid()* function.
  - 10011 • The [ELOOP] mandatory error condition is added.
  - 10012 • A second [ENAMETOOLONG] is added as an optional error condition.

- 10013       • The [ETXTBSY] optional error condition is added.
- 10014       The following changes were made to align with the IEEE P1003.1a draft standard:
- 10015       • The [EINVAL] mandatory error condition is added.
- 10016       • The [ELOOP] optional error condition is added.
- 10017       The description of CPU-time clock semantics is added for alignment with  
10018       IEEE Std. 1003.1d-1999.
- 10019       The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by adding semantics  
10020       for typed memory.
- 10021       The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |
- 10022       The description of tracing semantics is added for alignment with IEEE Std. 1003.1q-2000. |

## 10023 NAME

10024 exit, \_Exit, \_exit — terminate a process

## 10025 SYNOPSIS

10026 #include &lt;stdlib.h&gt;

10027 void exit(int status);

10028 #include &lt;unistd.h&gt;

10029 void \_Exit(int status);

10030 void \_exit(int status);

## 10031 DESCRIPTION

10032 CX The functionality described on this reference page for the *exit()* function is aligned with the  
 10033 ISO C standard. Any conflict between the requirements described here and the ISO C standard  
 10034 are unintentional. This volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

10035 The *exit()* function shall first call all functions registered by *atexit()*, in the reverse order of their  
 10036 registration, except that a function is called after any previously registered functions that had  
 10037 already been called at the time it was registered. Each function is called as many times as it was  
 10038 registered. If, during the call to any such function, a call to the *longjmp()* function is made that  
 10039 would terminate the call to the registered function, the behavior is undefined.

10040 If a function registered by a call to *atexit()* fails to return, the remaining registered functions shall  
 10041 not be called and the rest of the *exit()* processing shall not be completed. If *exit()* is called more  
 10042 than once, the effects are undefined.

10043 CX The *exit()* function then flushes all output streams, closes all open streams, and removes all files  
 10044 created by *tmpfile()*. Finally, control is returned to the host environment as described below. The  
 10045 values of *status* can be *EXIT\_SUCCESS* or *EXIT\_FAILURE*, as described in <stdlib.h>, or any  
 10046 CX implementation-defined value, although note that only the range 0 through 255 shall be  
 10047 available to a waiting parent process.

10048 The *\_Exit()* and *\_exit()* functions shall be functionally identical.

10049 CX The *\_Exit()*, *\_exit()*, and *exit()* functions shall terminate the calling process with the following  
 10050 consequences:

- 10051 XSI • All of the file descriptors, directory streams, conversion descriptors, and message catalog  
 10052 descriptors open in the calling process are closed.
- 10053 XSI • If the parent process of the calling process is executing a *wait()*, *waitid()*, or *waitpid()*, and has  
 10054 neither set its *SA\_NOCLDWAIT* flag nor set *SIGCHLD* to *SIG\_IGN*, it is notified of the  
 10055 calling process' termination and the low-order eight bits (that is, bits 0377) of *status* are made  
 10056 available to it. If the parent is not waiting, the child's status shall be made available to it  
 10057 XSI when the parent subsequently executes *wait()*, *waitid()*, or *waitpid()*.
- 10058 XSI • If the parent process of the calling process is not executing a *wait()*, *waitid()*, or *waitpid()*, and  
 10059 has not set its *SA\_NOCLDWAIT* flag, or set *SIGCHLD* to *SIG\_IGN*, the calling process is  
 10060 transformed into a *zombie process*. A *zombie process* is an inactive process and it shall be  
 10061 XSI deleted at some later time when its parent process executes *wait()*, *waitid()*, or *waitpid()*.
- 10062 • Termination of a process does not directly terminate its children. The sending of a *SIGHUP*  
 10063 signal as described below indirectly terminates children in some circumstances.
- 10064 • If the implementation supports the *SIGCHLD* signal, a *SIGCHLD* shall be sent to the parent  
 10065 process.

- 10066 XSI  
10067  
10068  
10069
- If the parent process has set its SA\_NOCLDWAIT flag, or set SIGCHLD to SIG\_IGN, the status shall be discarded, and the lifetime of the calling process shall end immediately. If SA\_NOCLDWAIT is set, it is implementation-defined whether a SIGCHLD signal shall be sent to the parent process.
- 10070  
10071  
10072
- The parent process ID of all of the calling process' existing child processes and zombie processes is set to the process ID of an implementation-defined system process. That is, these processes are inherited by a special system process.
- 10073 XSI  
10074
- Each attached shared-memory segment is detached and the value of *shm\_nattch* (see *shmget()*) in the data structure associated with its shared memory ID is decremented by 1.
- 10075 XSI  
10076
- For each semaphore for which the calling process has set a *semadj* value (see *semop()*), that value is added to the *semval* of the specified semaphore.
- 10077  
10078
- If the process is a controlling process, the SIGHUP signal shall be sent to each process in the foreground process group of the controlling terminal belonging to the calling process.
- 10079  
10080
- If the process is a controlling process, the controlling terminal associated with the session is disassociated from the session, allowing it to be acquired by a new controlling process.
- 10081  
10082  
10083
- If the exit of the process causes a process group to become orphaned, and if any member of the newly-orphaned process group is stopped, then a SIGHUP signal followed by a SIGCONT signal shall be sent to each process in the newly-orphaned process group.
- 10084 SEM  
10085
- All open named semaphores in the calling process shall be closed as if by appropriate calls to *sem\_close()*.
- 10086 ML  
10087  
10088  
10089
- Any memory locks established by the process via calls to *mlockall()* or *mlock()* shall be removed. If locked pages in the address space of the calling process are also mapped into the address spaces of other processes and are locked by those processes, the locks established by the other processes shall be unaffected by the call by this process to *\_Exit()* or *\_exit()*.
- 10090 MF|SHM
- Memory mappings created in the process are unmapped before the process is destroyed.
- 10091 TYM  
10092
- Any blocks of typed memory that were mapped in the calling process are unmapped, as if *munmap()* was implicitly called to unmap them.
- 10093 MSG  
10094
- All open message queue descriptors in the calling process shall be closed as if by appropriate calls to *mq\_close()*.
- 10095 AIO  
10096  
10097  
10098  
10099  
10100
- Any outstanding cancelable asynchronous I/O operations may be canceled. Those asynchronous I/O operations that are not canceled shall complete as if the *\_Exit()* or *\_exit()* operation had not yet occurred, but any associated signal notifications shall be suppressed. The *\_Exit()* or *\_exit()* operation may block awaiting such I/O completion. Whether any I/O is canceled, and which I/O may be canceled upon *\_Exit()* or *\_exit()*, is implementation-defined.
- 10101  
10102
- Threads terminated by a call to *\_Exit()* or *\_exit()* shall not invoke their cancelation cleanup handlers or per-thread data destructors.
- 10103 TRC  
10104  
10105  
10106
- If the calling process is a trace controller process, any trace streams that were created by the calling process shall be shut down as described by the *posix\_trace\_shutdown()* function, and any process' mapping of trace event names to trace event type identifiers built for these trace streams may be deallocated.

10107 **RETURN VALUE**

10108       These functions do not return.

10109 **ERRORS**

10110       No errors are defined.

10111 **EXAMPLES**

10112       None.

10113 **APPLICATION USAGE**

10114       Normally applications should use *exit()* rather than *\_Exit()* or *\_exit()*.

10115 **RATIONALE**10116       **Process Termination**

10117       Early proposals drew a different distinction between normal and abnormal process termination.  
10118       Abnormal termination was caused only by certain signals and resulted in implementation-  
10119       defined “actions”, as discussed below. Subsequent proposals distinguished three types of  
10120       termination: *normal termination* (as in the current specification), *simple abnormal termination*, and  
10121       *abnormal termination with actions*. Again the distinction between the two types of abnormal  
10122       termination was that they were caused by different signals and that implementation-defined  
10123       actions would result in the latter case. Given that these actions were completely  
10124       implementation-defined, the early proposals were only saying when the actions could occur and  
10125       how their occurrence could be detected, but not what they were. This was of little or no use to  
10126       portable applications, and thus the distinction is not made in this volume of  
10127       IEEE Std. 1003.1-200x.

10128       The implementation-defined actions usually include, in most historical implementations, the  
10129       creation of a file named **core** in the current working directory of the process. This file contains an  
10130       image of the memory of the process, together with descriptive information about the process,  
10131       perhaps sufficient to reconstruct the state of the process at the receipt of the signal.

10132       There is a potential security problem in creating a **core** file if the process was set-user-ID and the  
10133       current user is not the owner of the program, if the process was set-group-ID and none of the  
10134       user’s groups match the group of the program, or if the user does not have permission to write in  
10135       the current directory. In this situation, an implementation either should not create a **core** file or  
10136       should make it unreadable by the user.

10137       Despite the silence of this volume of IEEE Std. 1003.1-200x on this feature, applications are  
10138       advised not to create files named **core** because of potential conflicts in many implementations.  
10139       Some historical implementations use a different name than **core** for the file, such as by  
10140       appending the process ID to the file name.

10141       **Terminating a Process**

10142       It is important that the consequences of process termination as described occur regardless of  
10143       whether the process called *\_exit()* (perhaps indirectly through *exit()*) or instead was terminated  
10144       due to a signal or for some other reason. Note that in the specific case of *exit()* this means that  
10145       the *status* argument to *exit()* is treated the same as the *status* argument to *\_exit()*.

10146       A language other than C may have other termination primitives than the C-language *exit()*  
10147       function, and programs written in such a language should use its native termination primitives,  
10148       but those should have as part of their function the behavior of *\_exit()* as described.  
10149       Implementations in languages other than C are outside the scope of the present version of this  
10150       volume of IEEE Std. 1003.1-200x, however.

10151 As required by the ISO C standard, using **return** from *main()* is equivalent to calling *exit()* with  
10152 the same argument value. Also, reaching the end of the *main()* function is equivalent to using  
10153 *exit()* with an unspecified value.

10154 A value of zero (or `EXIT_SUCCESS`, which is required to be zero) for the argument *status*  
10155 conventionally indicates successful termination. This corresponds to the specification for *exit()*  
10156 in the ISO C standard. The convention is followed by utilities such as *make* and various shells,  
10157 which interpret a zero status from a child process as success. For this reason, applications should  
10158 not call *exit(0)* or *\_exit(0)* when they terminate unsuccessfully; for example, in signal-catching  
10159 functions.

10160 Historically, the implementation-defined process that inherits children whose parents have  
10161 terminated without waiting on them is called *init* and has a process ID of 1.

10162 The sending of a `SIGHUP` to the foreground process group when a controlling process  
10163 terminates corresponds to somewhat different historical implementations. In System V, the  
10164 kernel sends a `SIGHUP` on termination of (essentially) a controlling process. In 4.2 BSD, the  
10165 kernel does not send `SIGHUP` in a case like this, but the termination of a controlling process is  
10166 usually noticed by a system daemon, which arranges to send a `SIGHUP` to the foreground  
10167 process group with the *vhangup()* function. However, in 4.2 BSD, due to the behavior of the  
10168 shells that support job control, the controlling process is usually a shell with no other processes  
10169 in its process group. Thus, a change to make *\_exit()* behave this way in such systems should not  
10170 cause problems with existing applications.

10171 The termination of a process may cause a process group to become orphaned in either of two  
10172 ways. The connection of a process group to its parent(s) outside of the group depends on both  
10173 the parents and their children. Thus, a process group may be orphaned by the termination of the  
10174 last connecting parent process outside of the group or by the termination of the last direct  
10175 descendant of the parent process(es). In either case, if the termination of a process causes a  
10176 process group to become orphaned, processes within the group are disconnected from their job  
10177 control shell, which no longer has any information on the existence of the process group.  
10178 Stopped processes within the group would languish forever. In order to avoid this problem,  
10179 newly orphaned process groups that contain stopped processes are sent a `SIGHUP` signal and a  
10180 `SIGCONT` signal to indicate that they have been disconnected from their session. The `SIGHUP`  
10181 signal causes the process group members to terminate unless they are catching or ignoring  
10182 `SIGHUP`. Under most circumstances, all of the members of the process group are stopped if any  
10183 of them are stopped.

10184 The action of sending a `SIGHUP` and a `SIGCONT` signal to members of a newly orphaned  
10185 process group is similar to the action of 4.2 BSD, which sends `SIGHUP` and `SIGCONT` to each  
10186 stopped child of an exiting process. If such children exit in response to the `SIGHUP`, any  
10187 additional descendants receive similar treatment at that time. In this volume of  
10188 IEEE Std. 1003.1-200x, the signals are sent to the entire process group at the same time. Also, in  
10189 this volume of IEEE Std. 1003.1-200x, but not in 4.2 BSD, stopped processes may be orphaned,  
10190 but may be members of a process group that is not orphaned; therefore, the action taken at  
10191 *\_exit()* must consider processes other than child processes.

10192 It is possible for a process group to be orphaned by a call to *setpgid()* or *setsid()*, as well as by  
10193 process termination. This volume of IEEE Std. 1003.1-200x does not require sending `SIGHUP`  
10194 and `SIGCONT` in those cases, because, unlike process termination, those cases are not caused  
10195 accidentally by applications that are unaware of job control. An implementation can choose to  
10196 send `SIGHUP` and `SIGCONT` in those cases as an extension; such an extension must be  
10197 documented as required in `<signal.h>`.

10198 The ISO/IEC 9899:1999 standard adds the *\_Exit()* function that results in immediate program  
10199 termination without triggering signals or *atexit()*-registered functions. In IEEE Std. 1003.1-200x,

10200 this is equivalent to the `_exit()` function.

#### 10201 FUTURE DIRECTIONS

10202 None.

#### 10203 SEE ALSO

10204 `atexit()`, `close()`, `fclose()`, `longjmp()`, <REFERENCE UNDEFINED>(posix\_trace\_eventid),  
 10205 `posix_trace_shutdown()`, `posix_trace_trid_eventid_open()`, `semop()`, `shmget()`, `sigaction()`, `wait()`,  
 10206 `waitid()`, `waitpid()`, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>, <unistd.h>

#### 10207 CHANGE HISTORY

10208 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 10209 Issue 4

10210 The <unistd.h> header is added to the SYNOPSIS for `_exit()`.

10211 In the DESCRIPTION, text is added describing the behavior when a function registered by  
 10212 `atexit()` fails to return, and the consequences of calling `exit()` more than once.

10213 The phrase “If the implementation supports job control” is removed from the last bullet in the  
 10214 DESCRIPTION. This is because job control is now defined as mandatory for all conforming  
 10215 implementations.

10216 The following change is incorporated for alignment with the ISO C standard:

- 10217 • In the DESCRIPTION, interactions between `exit()` and `atexit()` are defined, and it is now
- 10218 stated explicitly that all files created by `tmpfile()` are removed.

#### 10219 Issue 4, Version 2

10220 The following changes to the DESCRIPTION are incorporated for X/OPEN UNIX conformance:

- 10221 • References to the functions `wait3()` and `waitid()` are added in appropriate places throughout
- 10222 the text.
- 10223 • Interactions with the SA\_NOCLDWAIT flag and SIGCHLD signal are defined.
- 10224 • It is specified that each mapped memory object is unmapped.

#### 10225 Issue 5

10226 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
 10227 Threads Extension.

10228 Interactions with the SA\_NOCLDWAIT flag and SIGCHLD signal are further clarified.

10229 The values of `status` from `exit()` are better described.

#### 10230 Issue 6

10231 Extensions beyond the ISO C standard are now marked.

10232 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by adding semantics  
 10233 for typed memory.

10234 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 10235 • The `_Exit()` function is included.
- 10236 • The DESCRIPTION is updated.

10237 The description of tracing semantics is added for alignment with IEEE Std. 1003.1q-2000.

10238 References to the `wait3()` function are removed.

10239 **NAME**

10240 exp, expf, expl — exponential function

10241 **SYNOPSIS**

10242 #include &lt;math.h&gt;

10243 double exp(double x);

10244 float expf(float x);

10245 long double expl(long double x);

10246 **DESCRIPTION**

10247 *CX* The functionality described on this reference page is aligned with the ISO C standard. Any  
10248 conflict between the requirements described here and the ISO C standard is unintentional. This  
10249 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

10250 These functions shall compute the exponent of  $x$ , defined as  $e^x$ .

10251 An application wishing to check for error situations should set *errno* to 0 before calling *exp()*. If  
10252 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

10253 **RETURN VALUE**10254 Upon successful completion, these functions shall return the exponential value of  $x$ .

10255 If the correct value would cause overflow, *exp()* shall return HUGE\_VAL and set *errno* to  
10256 [ERANGE].

10257 If the correct value would cause underflow, *exp()* shall return 0 and may set *errno* to [ERANGE].10258 *XSI* If  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].10259 **ERRORS**

10260 These functions shall fail if:

10261 [ERANGE] The result overflows.

10262 These functions may fail if:

10263 *XSI* [EDOM] The value of  $x$  is NaN.

10264 [ERANGE] The result underflows

10265 *XSI* No other errors shall occur.10266 **EXAMPLES**

10267 None.

10268 **APPLICATION USAGE**

10269 None.

10270 **RATIONALE**

10271 None.

10272 **FUTURE DIRECTIONS**

10273 None.

10274 **SEE ALSO**10275 *isnan()*, *log()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>10276 **CHANGE HISTORY**

10277 First released in Issue 1. Derived from Issue 1 of the SVID.

10278 **Issue 4**

10279       References to *matherr()* are removed.

10280       The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
10281       ISO C standard and to rationalize error handling in the mathematics functions.

10282       The return value specified for [EDOM] is marked as an extension. |

10283 **Issue 5**

10284       The DESCRIPTION is updated to indicate how an application should check for an error. This  
10285       text was previously published in the APPLICATION USAGE section. |

10286 **Issue 6**

10287       The *expf()* and *expl()* functions are added for alignment with the ISO/IEC 9899:1999 standard. |

10288 **NAME**

10289 exp2, exp2f, exp2l — exponential base 2 functions

10290 **SYNOPSIS**

10291 #include &lt;math.h&gt;

10292 double exp2(double x);

10293 float exp2f(float x);

10294 long double exp2l(long double x);

10295 **DESCRIPTION**

10296 **cx** The functionality described on this reference page is aligned with the ISO C standard. Any  
10297 conflict between the requirements described here and the ISO C standard is unintentional. This  
10298 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

10299 These functions shall compute the base 2 exponent of  $x$ , defined as  $e^x$ .

10300 An application wishing to check for error situations should set *errno* to 0 before calling these  
10301 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

10302 **RETURN VALUE**10303 Upon successful completion, these functions shall return  $2^x$ .

10304 If the correct value would cause overflow, these functions shall return HUGE\_VAL and set *errno*  
10305 to [ERANGE].

10306 If the correct value would cause underflow, these functions shall return 0 and may set *errno* to  
10307 [ERANGE].

10308 If  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].10309 **ERRORS**

10310 These functions shall fail if:

10311 [ERANGE] The result overflows.

10312 These functions may fail if:

10313 [EDOM] The value of  $x$  is NaN.

10314 [ERANGE] The result underflows

10315 No other errors shall occur.

10316 **EXAMPLES**

10317 None.

10318 **APPLICATION USAGE**

10319 None.

10320 **RATIONALE**

10321 None.

10322 **FUTURE DIRECTIONS**

10323 None.

10324 **SEE ALSO**10325 *exp()*, *isnan()*, *log()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

10326 **CHANGE HISTORY**

10327 First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard.

10328 **NAME**

10329 expm1, expm1f, expm1l — compute exponential functions

10330 **SYNOPSIS**

```
10331 #include <math.h>
10332 double expm1(double x);
10333 float expm1f(float x);
10334 long double expm1l(long double x);
```

10335 **DESCRIPTION**

10336 These functions shall compute  $e^x - 1.0$ .

10337 **RETURN VALUE**

10338 If  $x$  is NaN, then these functions shall return NaN and *errno* may be set to [EDOM].

10339 If  $x$  is positive infinity, these functions shall return positive infinity.

10340 If  $x$  is negative infinity, these functions shall return  $-1.0$ .

10341 If the value overflows, these functions shall return HUGE\_VAL and may set *errno* to [ERANGE].

10342 **ERRORS**

10343 These functions may fail if:

10344 [EDOM] The value of  $x$  is NaN.

10345 [ERANGE] The result overflows.

10346 **EXAMPLES**

10347 None.

10348 **APPLICATION USAGE**

10349 The value of *expm1*( $x$ ) may be more accurate than *exp*( $x$ ) $-1.0$  for small values of  $x$ .

10350 The *expm1*() and *log1p*() functions are useful for financial calculations of  $((1+x)^n - 1)/x$ , namely:

```
10351 expm1(n * log1p(x))/x
```

10352 when  $x$  is very small (for example, when calculating small daily interest rates). These functions  
10353 also simplify writing accurate inverse hyperbolic functions.

10354 **RATIONALE**

10355 None.

10356 **FUTURE DIRECTIONS**

10357 None.

10358 **SEE ALSO**

10359 *exp*(), *ilogb*(), *log1p*(), the Base Definitions volume of IEEE Std. 1003.1-200x, **<math.h>**

10360 **CHANGE HISTORY**

10361 First released in Issue 4, Version 2.

10362 **Issue 5**

10363 Moved from X/OPEN UNIX extension to BASE.

10364 **Issue 6**

10365 The *expm1f*() and *expm1l*() functions are added for alignment with the ISO/IEC 9899:1999  
10366 standard.

10367 **NAME**

10368 fabs, fabsf, fabsl — absolute value function

10369 **SYNOPSIS**

10370 #include &lt;math.h&gt;

10371 double fabs(double x);

10372 float fabsf(float x);

10373 long double fabsl(long double x);

10374 **DESCRIPTION**

10375 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 10376 conflict between the requirements described here and the ISO C standard is unintentional. This  
 10377 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

10378 These functions shall compute the absolute value of  $x$ ,  $|x|$ .

10379 An application wishing to check for error situations should set *errno* to 0 before calling *fabs()*. If  
 10380 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

10381 **RETURN VALUE**10382 Upon successful completion, these functions shall return the absolute value of  $x$ .10383 **XSI** If  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].10384 If the result underflows, 0 shall be returned and *errno* may be set to [ERANGE].10385 **ERRORS**

10386 These functions may fail if:

10387 **XSI** [EDOM] The value of  $x$  is NaN.

10388 [ERANGE] The result underflows

10389 **XSI** No other errors shall occur.10390 **EXAMPLES**

10391 None.

10392 **APPLICATION USAGE**

10393 None.

10394 **RATIONALE**

10395 None.

10396 **FUTURE DIRECTIONS**

10397 None.

10398 **SEE ALSO**10399 *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>10400 **CHANGE HISTORY**

10401 First released in Issue 1. Derived from Issue 1 of the SVID.

10402 **Issue 4**10403 References to *matherr()* are removed.

10404 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
 10405 ISO C standard and to rationalize error handling in the mathematics functions.

10406 The return value specified for [EDOM] is marked as an extension.

10407 **Issue 5**

10408

10409

The DESCRIPTION is updated to indicate how an application should check for an error. This text was previously published in the APPLICATION USAGE section.

10410 **Issue 6**

10411

The *fabsf()* and *fabsl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

10412 **NAME**

10413 fattach — attach a STREAMS-based file descriptor to a file in the file system name space  
 10414 (**STREAMS**)

10415 **SYNOPSIS**

```
10416 XSR #include <stropts.h>
10417 int fattach(int fildev, const char *path);
10418
```

10419 **DESCRIPTION**10420 **Notes to Reviewers**

10421 *This section with side shading will not appear in the final copy. - Ed.*

10422 Re D1, XSH, ERN 111: if the original file had multiple links, the streams file still has only one? I  
 10423 presume that a stream is actually attached to an inode, not a file name. If so, there continue to  
 10424 exist multiple links to the object, even though it shows a link count of 1. If it associated the  
 10425 stream with a file name, then the following sentence is wrong.

10426 The *fattach()* function attaches a STREAMS-based file descriptor to a file, effectively associating  
 10427 a path name with *fildev*. The application shall ensure that the *fildev* argument is a valid open file  
 10428 descriptor associated with a STREAMS file. The *path* argument points to a path name of an  
 10429 existing file. The application shall ensure that the process has appropriate privileges, or is the  
 10430 owner of the file named by *path* and has write permission. A successful call to *fattach()* shall  
 10431 cause all path names that name the file named by *path* to name the STREAMS file associated  
 10432 with *fildev*, until the STREAMS file is detached from the file. A STREAMS file can be attached to  
 10433 more than one file and can have several path names associated with it.

10434 The attributes of the named STREAMS file shall be initialized as follows: the permissions, user  
 10435 ID, group ID, and times are set to those of the file named by *path*, the number of links is set to 1,  
 10436 and the size and device identifier are set to those of the STREAMS file associated with *fildev*. If  
 10437 any attributes of the named STREAMS file are subsequently changed (for example, by *chmod()*),  
 10438 neither the attributes of the underlying file nor the attributes of the STREAMS file to which *fildev*  
 10439 refers shall be affected.

10440 File descriptors referring to the underlying file, opened prior to an *fattach()* call, shall continue to  
 10441 refer to the underlying file.

10442 **RETURN VALUE**

10443 Upon successful completion, *fattach()* shall return 0. Otherwise, -1 shall be returned and *errno*  
 10444 set to indicate the error.

10445 **ERRORS**

10446 The *fattach()* function shall fail if:

10447 [EACCES] Search permission is denied for a component of the path prefix, or the process  
 10448 is the owner of *path* but does not have write permissions on the file named by  
 10449 *path*.

10450 [EBADF] The *fildev* argument is not a valid open file descriptor.

10451 [EBUSY] The file named by *path* is currently a mount point or has a STREAMS file  
 10452 attached to it.

10453 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 10454 argument.

|       |                |                                                                                              |
|-------|----------------|----------------------------------------------------------------------------------------------|
| 10455 | [ENAMETOOLONG] |                                                                                              |
| 10456 |                | The size of <i>path</i> exceeds {PATH_MAX} or a component of <i>path</i> is longer than      |
| 10457 |                | {NAME_MAX}.                                                                                  |
| 10458 | [ENOENT]       | A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string. |
| 10459 | [ENOTDIR]      | A component of the path prefix is not a directory.                                           |
| 10460 | [EPERM]        | The effective user ID of the process is not the owner of the file named by <i>path</i>       |
| 10461 |                | and the process does not have appropriate privilege.                                         |
| 10462 |                | The <i>fattach()</i> function may fail if:                                                   |
| 10463 | [EINVAL]       | The <i>fildev</i> argument does not refer to a STREAMS file.                                 |
| 10464 | [ELOOP]        | More than {SYMLOOP_MAX} symbolic links were encountered during                               |
| 10465 |                | resolution of the <i>path</i> argument.                                                      |
| 10466 | [ENAMETOOLONG] |                                                                                              |
| 10467 |                | Path name resolution of a symbolic link produced an intermediate result                      |
| 10468 |                | whose length exceeds {PATH_MAX}.                                                             |
| 10469 | [EXDEV]        | A link to a file on another file system was attempted.                                       |

#### 10470 EXAMPLES

##### 10471 Attaching a File Descriptor to a File

10472 In the following example, *fd* refers to an open STREAMS file. The call to *fattach()* associates this  
 10473 STREAM with the file **/tmp/named-STREAM**, such that any future calls to open **/tmp/named-**  
 10474 **STREAM**, prior to breaking the attachment via a call to *fdetach()*, will instead create a new file  
 10475 handle referring to the STREAMS file associated with *fd*.

```
10476 #include <stropts.h>
10477 ...
10478 int fd;
10479 char *filename = "/tmp/named-STREAM";
10480 int ret;
10481 ret = fattach(fd, filename);
```

#### 10482 APPLICATION USAGE

10483 The *fattach()* function behaves similarly to the traditional *mount()* function in the way a file is  
 10484 temporarily replaced by the root directory of the mounted file system. In the case of *fattach()*, the  
 10485 replaced file need not be a directory and the replacing file is a STREAMS file.

#### 10486 RATIONALE

10487 None.

#### 10488 FUTURE DIRECTIONS

10489 None.

#### 10490 SEE ALSO

10491 *fdetach()*, *isastream()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<stropts.h>**

#### 10492 CHANGE HISTORY

10493 First released in Issue 4, Version 2.

10494 **Issue 5**

10495 Moved from X/OPEN UNIX extension to BASE.

10496 The [EXDEV] error is added to the list of optional errors in the ERRORS section.

10497 **Issue 6**

10498 This function is marked as part of the XSI STREAMS Option Group.

10499 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

10500 The wording of the mandatory [ELOOP] error condition is updated, and a second optional |  
10501 [ELOOP] error condition is added. |

10502 **NAME**

10503 fchdir — change working directory

10504 **SYNOPSIS**

10505 XSI #include &lt;unistd.h&gt;

10506 int fchdir(int *fildev*);

10507

10508 **DESCRIPTION**10509 The *fchdir()* function has the same effect as *chdir()* except that the directory that is to be the new  
10510 current working directory is specified by the file descriptor *fildev*.10511 **RETURN VALUE**10512 Upon successful completion, *fchdir()* shall return 0. Otherwise, it shall return -1 and set *errno* to  
10513 indicate the error. On failure the current working directory shall remain unchanged.10514 **ERRORS**10515 The *fchdir()* function shall fail if:10516 [EACCES] Search permission is denied for the directory referenced by *fildev*.10517 [EBADF] The *fildev* argument is not an open file descriptor.10518 [ENOTDIR] The open file descriptor *fildev* does not refer to a directory.10519 The *fchdir()* may fail if:10520 [EINTR] A signal was caught during the execution of *fchdir()*.

10521 [EIO] An I/O error occurred while reading from or writing to the file system.

10522 **EXAMPLES**

10523 None.

10524 **APPLICATION USAGE**

10525 None.

10526 **RATIONALE**

10527 None.

10528 **FUTURE DIRECTIONS**

10529 None.

10530 **SEE ALSO**10531 *chdir()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <unistd.h>10532 **CHANGE HISTORY**

10533 First released in Issue 4, Version 2.

10534 **Issue 5**

10535 Moved from X/OPEN UNIX extension to BASE.

10536 **NAME**

10537 fchmod — change mode of a file

10538 **SYNOPSIS**

10539 #include &lt;sys/stat.h&gt;

10540 int fchmod(int *fildev*, mode\_t *mode*);10541 **DESCRIPTION**10542 The *fchmod()* function has the same effect as *chmod()* except that the file whose permissions are  
10543 changed is specified by the file descriptor *fildev*.10544 SHM If *fildev* references a shared memory object, the *fchmod()* function need only affect the S\_IRUSR,  
10545 S\_IWUSR, S\_IRGRP, S\_IWGRP, S\_IROTH, and S\_IWOTH file permission bits.10546 TYM If *fildev* references a typed memory object, the behavior of *fchmod()* is unspecified.10547 **Notes to Reviewers**10548 *This section with side shading will not appear in the final copy. - Ed.*10549 D3, XSH, ERN 178 suggests adding text as follows: "If *fildev* refers to a STREAM (which is  
10550 fattached() into the file system name space) the call returns successfully, doing nothing. If *fildev*  
10551 refers to a stream, <do what?>."10552 **RETURN VALUE**10553 Upon successful completion, *fchmod()* shall return 0. Otherwise, it shall return -1 and set *errno* to  
10554 indicate the error.10555 **ERRORS**10556 The *fchmod()* function shall fail if:10557 [EBADF] The *fildev* argument is not an open file descriptor.10558 [EPERM] The effective user ID does not match the owner of the file and the process  
10559 does not have appropriate privilege.10560 [EROFS] The file referred to by *fildev* resides on a read-only file system.10561 The *fchmod()* function may fail if:10562 XSI [EINTR] The *fchmod()* function was interrupted by a signal.10563 XSI [EINVAL] The value of the *mode* argument is invalid.10564 [EINVAL] The *fildev* argument refers to a pipe and the implementation disallows  
10565 execution of *fchmod()* on a pipe.10566 **EXAMPLES**10567 **Changing the Current Permissions for a File**10568 The following example shows how to change the permissions for a file named */home/cnd/mod1*  
10569 so that the owner and group have read/write/execute permissions, but the world only has  
10570 read/write permissions.

10571 #include &lt;sys/stat.h&gt;

10572 #include &lt;fcntl.h&gt;

10573 mode\_t mode;

10574 int fildev;

10575 ...

```
10576 fildes = open("/home/cnd/mod1", O_RDWR);
10577 fchmod(fildes, S_IRWXU | S_IRWXG | S_IROTH | S_IWOTH);
```

**10578 APPLICATION USAGE**

10579 None.

**10580 RATIONALE**

10581 None.

**10582 FUTURE DIRECTIONS**

10583 None.

**10584 SEE ALSO**

10585 *chmod()*, *chown()*, *creat()*, *fcntl()*, *fstatvfs()*, *mknod()*, *open()*, *read()*, *stat()*, *write()*, the Base  
10586 Definitions volume of IEEE Std. 1003.1-200x, <sys/stat.h>

**10587 CHANGE HISTORY**

10588 First released in Issue 4, Version 2.

**10589 Issue 5**

10590 Moved from X/OPEN UNIX extension to BASE and aligned with *fchmod()* in the POSIX  
10591 Realtime Extension. Specifically, the second paragraph of the DESCRIPTION is added and a  
10592 second instance of [EINVAL] is defined in the list of optional errors.

**10593 Issue 6**

10594 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by stating that  
10595 *fchmod()* behavior is unspecified for typed memory objects.

10596 **NAME**

10597 fchown — change owner and group of a file

10598 **SYNOPSIS**

10599 #include &lt;unistd.h&gt;

10600 int fchown(int *fildev*, uid\_t *owner*, gid\_t *group*);10601 **DESCRIPTION**10602 The *fchown()* function has the same effect as *chown()* except that the file whose owner and group  
10603 are changed is specified by the file descriptor *fildev*.10604 **RETURN VALUE**10605 Upon successful completion, *fchown()* shall return 0. Otherwise, it shall return -1 and set *errno* to  
10606 indicate the error.10607 **ERRORS**10608 The *fchown()* function shall fail if:10609 [EBADF] The *fildev* argument is not an open file descriptor.10610 [EPERM] The effective user ID does not match the owner of the file or the process does  
10611 not have appropriate privilege and `_POSIX_CHOWN_RESTRICTED` indicates  
10612 that such privilege is required.10613 [EROFS] The file referred to by *fildev* resides on a read-only file system.10614 The *fchown()* function may fail if:10615 [EINVAL] The owner or group ID is not a value supported by the implementation. The  
10616 *fildev* argument refers to a pipe and the implementation disallows execution of  
10617 *fchown()* on a pipe.10618 **Notes to Reviewers**10619 *This section with side shading will not appear in the final copy. - Ed.*10620 D3, XSH, ERN 177 states that STREAMS ignore the call, but raises a question  
10621 about AF\_UNIX sockets in the file system name space.

10622 [EIO] A physical I/O error has occurred.

10623 [EINTR] The *fchown()* function was interrupted by a signal which was caught.10624 **EXAMPLES**10625 **Changing the Current Owner of a File**10626 The following example shows how to change the owner of a file named `/home/cnd/mod1` to  
10627 “jones” and the group to “cnd”.10628 The numeric value for the user ID is obtained by extracting the user ID from the user database  
10629 entry associated with “jones”. Similarly, the numeric value for the group ID is obtained by  
10630 extracting the group ID from the group database entry associated with “cnd”. This example  
10631 assumes the calling program has appropriate privileges.

10632 #include &lt;sys/types.h&gt;

10633 #include &lt;unistd.h&gt;

10634 #include &lt;fcntl.h&gt;

10635 #include &lt;pwd.h&gt;

10636 #include &lt;grp.h&gt;

```
10637 struct passwd *pwd;
10638 struct group *grp;
10639 int fildes;
10640 ...
10641 fildes = open("/home/cnd/mod1", O_RDWR);
10642 pwd = getpwnam("jones");
10643 grp = getgrnam("cnd");
10644 fchown(fildes, pwd->pw_uid, grp->gr_gid);
```

**10645 APPLICATION USAGE**

10646 None.

**10647 RATIONALE**

10648 None.

**10649 FUTURE DIRECTIONS**

10650 None.

**10651 SEE ALSO**

10652 *chown()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

**10653 CHANGE HISTORY**

10654 First released in Issue 4, Version 2.

**10655 Issue 5**

10656 Moved from X/OPEN UNIX extension to BASE.

**10657 Issue 6**

10658 The following changes were made to align with the IEEE P1003.1a draft standard:

10659 • Clarification is added that a call to *fchown()* may not be allowed on a pipe.

10660 The *fchwon()* function is now defined as mandatory.

10661 **NAME**

10662        fclose — close a stream

10663 **SYNOPSIS**

10664        #include &lt;stdio.h&gt;

10665        int fclose(FILE \*stream);

10666 **DESCRIPTION**

10667 CX        The functionality described on this reference page is aligned with the ISO C standard. Any  
 10668        conflict between the requirements described here and the ISO C standard is unintentional. This  
 10669        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

10670        The *fclose()* function shall cause the stream pointed to by *stream* to be flushed and the associated  
 10671        file to be closed. Any unwritten buffered data for the stream shall be written to the file; any  
 10672        unread buffered data shall be discarded. Whether or not the call succeeds, the stream shall be  
 10673        disassociated from the file and any buffer set by the *setbuf()* or *setvbuf()* function shall be  
 10674        disassociated from the stream. If the associated buffer was automatically allocated, it shall be  
 10675 CX        deallocated. It shall mark for update the *st\_ctime* and *st\_mtime* fields of the underlying file, if the  
 10676        stream was writable, and if buffered data remains that has not yet been written to the file. The  
 10677        *fclose()* function shall perform the equivalent of a *close()* on the file descriptor that is associated  
 10678        with the stream pointed to by *stream*.

10679        After the call to *fclose()*, any use of *stream* results in undefined behavior.

10680 **RETURN VALUE**

10681 CX        Upon successful completion, *fclose()* shall return 0; otherwise, it shall return EOF and set *errno* to  
 10682        indicate the error.

10683 **ERRORS**

10684        The *fclose()* function shall fail if:

10685 CX        [EAGAIN]        The O\_NONBLOCK flag is set for the file descriptor underlying *stream* and the  
 10686        process would be delayed in the write operation.

10687 CX        [EBADF]        The file descriptor underlying stream is not valid.

10688 CX        [EFBIG]        An attempt was made to write a file that exceeds the maximum file size.

10689 XSI       [EFBIG]        An attempt was made to write a file that exceeds the process' file size limit.

10690 CX        [EFBIG]        The file is a regular file and an attempt was made to write at or beyond the  
 10691        offset maximum associated with the corresponding stream.

10692 CX        [EINTR]        The *fclose()* function was interrupted by a signal.

10693 CX        [EIO]            The process is a member of a background process group attempting to write  
 10694        to its controlling terminal, TOSTOP is set, the process is neither ignoring nor  
 10695        blocking SIGTTOU, and the process group of the process is orphaned. This  
 10696        error may also be returned under implementation-defined conditions.

10697 CX        [ENOSPC]        There was no free space remaining on the device containing the file.

10698 CX        [EPIPE]        An attempt is made to write to a pipe or FIFO that is not open for reading by  
 10699        any process. A SIGPIPE signal shall also be sent to the thread.

10700        The *fclose()* function may fail if:

10701 CX        [ENXIO]        A request was made of a nonexistent device, or the request was outside the  
 10702        capabilities of the device.

10703 **EXAMPLES**

10704 None.

10705 **APPLICATION USAGE**

10706 None.

10707 **RATIONALE**

10708 None.

10709 **FUTURE DIRECTIONS**

10710 None.

10711 **SEE ALSO**

10712 *close()*, *fopen()*, *getrlimit()*, *ulimit()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
10713 `<stdio.h>`

10714 **CHANGE HISTORY**

10715 First released in Issue 1. Derived from Issue 1 of the SVID.

10716 **Issue 4**

10717 The last sentence of the first paragraph in the DESCRIPTION is changed to say *close()* instead of  
10718 *fclose()*. This was an error in Issue 3.

10719 The following paragraph is withdrawn from the DESCRIPTION (by POSIX as well as X/Open)  
10720 because of the possibility of causing applications to malfunction, and the impossibility of  
10721 implementing these mechanisms for pipes:

10722 If the file is not already at EOF, and the file is one capable of seeking, the file *offset* of the  
10723 underlying open file description is adjusted so that the next operation on the open file  
10724 description deals with the byte after the last one read from or written to the stream being  
10725 closed.

10726 It is replaced with a statement that any subsequent use of *stream* is undefined.

10727 The [EFBIG] error is marked to indicate the extensions.

10728 **Issue 4, Version 2**10729 A cross-reference to *getrlimit()* is added.10730 **Issue 5**

10731 Large File Summit extensions are added.

10732 **Issue 6**

10733 Extensions beyond the ISO C standard are now marked.

10734 The following new requirements on POSIX implementations derive from alignment with the  
10735 Single UNIX Specification:

10736 • The [EFBIG] error is added as part of the large file support extensions.

10737 • The [ENXIO] optional error condition is added.

10738 The DESCRIPTION is updated to note that the stream and any buffer are disassociated whether  
10739 or not the call succeeds. This is for alignment with the ISO/IEC 9899:1999 standard.

## 10740 NAME

10741 fcntl — file control

## 10742 SYNOPSIS

10743 OH #include &lt;unistd.h&gt;

10744 #include &lt;fcntl.h&gt;

10745 int fcntl(int *fildes*, int *cmd*, ...);

## 10746 DESCRIPTION

10747 The *fcntl()* function provides for control over open files. The *fildes* argument is a file descriptor.10748 The available values for *cmd* are defined in <fcntl.h>, which include:

10749 **F\_DUPFD** Return a new file descriptor which is the lowest numbered available (that is,  
10750 not already open) file descriptor greater than or equal to the third argument,  
10751 *arg*, taken as an integer of type **int**. The new file descriptor refers to the same  
10752 open file description as the original file descriptor, and shares any locks. The  
10753 **FD\_CLOEXEC** flag associated with the new file descriptor is cleared to keep  
10754 the file open across calls to one of the *exec* functions.

10755 **F\_GETFD** Get the file descriptor flags defined in <fcntl.h> that are associated with the  
10756 file descriptor *fildes*. File descriptor flags are associated with a single file  
10757 descriptor and do not affect other file descriptors that refer to the same file.

10758 **F\_SETFD** Set the file descriptor flags defined in <fcntl.h>, that are associated with *fildes*,  
10759 to the third argument, *arg*, taken as type **int**. If the **FD\_CLOEXEC** flag in the  
10760 third argument is 0, the file shall remain open across the *exec* functions;  
10761 otherwise, the file shall be closed upon successful execution of one of the *exec*  
10762 functions.

10763 **F\_GETFL** Get the file status flags and file access modes, defined in <fcntl.h>, for the file  
10764 description associated with *fildes*. The file access modes can be extracted from  
10765 the return value using the mask **O\_ACCMODE**, which is defined in <fcntl.h>.  
10766 File status flags and file access modes are associated with the file description  
10767 and do not affect other file descriptors that refer to the same file with different  
10768 open file descriptions.

10769 **F\_SETFL** Set the file status flags, defined in <fcntl.h>, for the file description associated  
10770 with *fildes* from the corresponding bits in the third argument, *arg*, taken as  
10771 type **int**. Bits corresponding to the file access mode and the *oflag* values that  
10772 are set in *arg* are ignored. If any bits in *arg* other than those mentioned here are  
10773 changed by the application, the result is unspecified.

10774 **F\_GETOWN** If *fildes* refers to a socket, get the process or process group ID specified to  
10775 receive **SIGURG** signals when out-of-band data is available. Positive values  
10776 indicate a process ID; negative values, other than -1, indicate a process group  
10777 ID. If *fildes* does not refer to a socket, the results are unspecified.

10778 **F\_SETOWN** If *fildes* refers to a socket, set the process or process group ID specified to  
10779 receive **SIGURG** signals when out-of-band data is available, using the value of  
10780 the third argument, *arg*, taken as type **int**. Positive values indicate a process  
10781 ID; negative values, other than -1, indicate a process group ID. If *fildes* does  
10782 not refer to a socket, the results are unspecified.

10783 The following values for *cmd* are available for advisory record locking. Record locking is  
10784 supported for regular files, and may be supported for other files.

10785           F\_GETLK           Get the first lock which blocks the lock description pointed to by the third  
10786                           argument, *arg*, taken as a pointer to type **struct flock**, defined in `<fcntl.h>`.  
10787                           The information retrieved overwrites the information passed to *fcntl()* in the  
10788                           structure **flock**. If no lock is found that would prevent this lock from being  
10789                           created, then the structure shall be left unchanged except for the lock type  
10790                           which shall be set to F\_UNLCK.

10791           F\_SETLK           Set or clear a file segment lock according to the lock description pointed to by  
10792                           the third argument, *arg*, taken as a pointer to type **struct flock**, defined in  
10793                           `<fcntl.h>`. F\_SETLK is used to establish shared (or read) locks (F\_RDLCK) or  
10794                           exclusive (or write) locks (F\_WRLCK), as well as to remove either type of lock  
10795                           (F\_UNLCK). F\_RDLCK, F\_WRLCK, and F\_UNLCK are defined in `<fcntl.h>`.  
10796                           If a shared or exclusive lock cannot be set, *fcntl()* shall return immediately  
10797                           with a return value of -1.

10798           F\_SETLKW          This command is the same as F\_SETLK except that if a shared or exclusive  
10799                           lock is blocked by other locks, the thread shall wait until the request can be  
10800                           satisfied. If a signal that is to be caught is received while *fcntl()* is waiting for a  
10801                           region, *fcntl()* shall be interrupted. Upon return from the signal handler,  
10802                           *fcntl()* shall return -1 with *errno* set to [EINTR], and the lock operation shall  
10803                           not be done.

10804           Additional implementation-defined values for *cmd* may be defined in `<fcntl.h>`. Their names  
10805           shall start with F\_.

10806           When a shared lock is set on a segment of a file, other processes shall be able to set shared locks  
10807           on that segment or a portion of it. A shared lock prevents any other process from setting an  
10808           exclusive lock on any portion of the protected area. A request for a shared lock shall fail if the  
10809           file descriptor was not opened with read access.

10810           An exclusive lock shall prevent any other process from setting a shared lock or an exclusive lock  
10811           on any portion of the protected area. A request for an exclusive lock shall fail if the file  
10812           descriptor was not opened with write access.

10813           The structure **flock** describes the type (*l\_type*), starting offset (*l\_whence*), relative offset (*l\_start*),  
10814           size (*l\_len*), and process ID (*l\_pid*) of the segment of the file to be affected.

10815           The value of *l\_whence* is {SEEK\_SET}, {SEEK\_CUR}, or {SEEK\_END}, to indicate that the relative  
10816           offset *l\_start* bytes shall be measured from the start of the file, current position, or end of the file,  
10817           respectively. The value of *l\_len* is the number of consecutive bytes to be locked. The value of  
10818           *l\_len* may be negative (where the definition of **off\_t** permits negative values of *l\_len*). The *l\_pid*  
10819           field is only used with F\_GETLK to return the process ID of the process holding a blocking lock.  
10820           After a successful F\_GETLK request, when a blocking lock is found, the values returned in the  
10821           **flock** structure shall be as follows:

10822           *l\_type*          Type of blocking lock found.

10823           *l\_whence*        {SEEK\_SET}.

10824           *l\_start*         Start of the blocking lock.

10825           *l\_len*          Length of the blocking lock.

10826           *l\_pid*          Process ID of the process that holds the blocking lock.

10827           If the command is F\_SETLKW and the process must wait for another process to release a lock,  
10828           then the range of bytes to be locked shall be determined before the *fcntl()* function blocks. If the  
10829           file size or file descriptor seek offset change while *fcntl()* is blocked, this shall not affect the  
10830           range of bytes locked.

10831 If *l\_len* is positive, the area affected starts at *l\_start* and ends at *l\_start*+ *l\_len*-1. If *l\_len* is  
 10832 negative, the area affected starts at *l\_start*+ *l\_len* and ends at *l\_start*-1. Locks may start and  
 10833 extend beyond the current end of a file, but the application shall ensure that they are not  
 10834 negative relative to the beginning of the file. A lock shall be set to extend to the largest possible  
 10835 value of the file offset for that file by setting *l\_len* to 0. If such a lock also has *l\_start* set to 0 and  
 10836 *l\_whence* is set to {SEEK\_SET}, the whole file shall be locked.

10837 There shall be at most one type of lock set for each byte in the file. Before a successful return  
 10838 from an F\_SETLK or an F\_SETLKW request when the calling process has previously existing  
 10839 locks on bytes in the region specified by the request, the previous lock type for each byte in the  
 10840 specified region shall be replaced by the new lock type. As specified above under the  
 10841 descriptions of shared locks and exclusive locks, an F\_SETLK or an F\_SETLKW request  
 10842 (respectively) shall fail or block when another process has existing locks on bytes in the specified  
 10843 region and the type of any of those locks conflicts with the type specified in the request.

10844 All locks associated with a file for a given process shall be removed when a file descriptor for  
 10845 that file is closed by that process or the process holding that file descriptor terminates. Locks are  
 10846 not inherited by a child process.

10847 A potential for deadlock occurs if a process controlling a locked region is put to sleep by  
 10848 attempting to lock another process' locked region. If the system detects that sleeping until a  
 10849 locked region is unlocked would cause a deadlock, *fcntl()* shall fail with an [EDEADLK] error.

10850 SHM When the file descriptor *fildev* refers to a shared memory object, the behavior of *fcntl()* shall be  
 10851 the same as for a regular file except the effect of the following values for the argument *cmd* shall  
 10852 be unspecified: F\_SETFL, F\_GETLK, F\_SETLK, and F\_SETLKW.

10853 TYM If *fildev* refers to a typed memory object, the result of the *fcntl()* function is unspecified.

10854 An unlock (F\_UNLCK) request in which *l\_len* is non-zero and the offset of the last byte of the  
 10855 requested segment is the maximum value for an object of type *off\_t*, when the process has an  
 10856 existing lock in which *l\_len* is 0 and which includes the last byte of the requested segment, shall  
 10857 be treated as a request to unlock from the start of the requested segment with an *l\_len* equal to 0.  
 10858 Otherwise, an unlock (F\_UNLCK) request shall attempt to unlock only the requested segment.

#### 10859 RETURN VALUE

10860 Upon successful completion, the value returned shall depend on *cmd* as follows:

|       |          |                                                                                |
|-------|----------|--------------------------------------------------------------------------------|
| 10861 | F_DUPFD  | A new file descriptor.                                                         |
| 10862 | F_GETFD  | Value of flags defined in <fcntl.h>. The return value shall not be negative.   |
| 10863 | F_SETFD  | Value other than -1.                                                           |
| 10864 | F_GETFL  | Value of file status flags and access modes. The return value is not negative. |
| 10865 | F_SETFL  | Value other than -1.                                                           |
| 10866 | F_GETLK  | Value other than -1.                                                           |
| 10867 | F_SETLK  | Value other than -1.                                                           |
| 10868 | F_SETLKW | Value other than -1.                                                           |
| 10869 | F_GETOWN | Value of the socket owner process or process group; this will not be -1.       |
| 10870 | F_SETOWN | Value other than -1.                                                           |
| 10871 |          | Otherwise, -1 shall be returned and <i>errno</i> set to indicate the error.    |

10872 **ERRORS**10873 The *fcntl()* function shall fail if:

10874 [EACCES] or [EAGAIN]

10875 The *cmd* argument is F\_SETLK; the type of lock (*l\_type*) is a shared (F\_RDLCK)  
 10876 or exclusive (F\_WRLCK) lock and the segment of a file to be locked is already  
 10877 exclusive-locked by another process, or the type is an exclusive lock and some  
 10878 portion of the segment of a file to be locked is already shared-locked or  
 10879 exclusive-locked by another process.

10880 [EBADF]

10881 The *fildev* argument is not a valid open file descriptor, or the argument *cmd* is  
 10882 F\_SETLK or F\_SETLKW, the type of lock, *l\_type*, is a shared lock (F\_RDLCK),  
 10883 and *fildev* is not a valid file descriptor open for reading, or the type of lock  
 10884 *l\_type*, is an exclusive lock (F\_WRLCK), and *fildev* is not a valid file descriptor  
 open for writing.

10885 [EINTR]

10885 The *cmd* argument is F\_SETLKW and the function was interrupted by a signal.

10886 [EINVAL]

10887 The *cmd* argument is invalid, or the *cmd* argument is F\_DUPFD and *arg* is  
 10888 negative or greater than or equal to {OPEN\_MAX}, or the *cmd* argument is  
 10889 F\_GETLK, F\_SETLK, or F\_SETLKW and the data pointed to by *arg* is not valid,  
 or *fildev* refers to a file that does not support locking.

10890 [EMFILE]

10891 The argument *cmd* is F\_DUPFD and {OPEN\_MAX} file descriptors are  
 10892 currently open in the calling process, or no file descriptors greater than or  
 equal to *arg* are available.

10893 [ENOLCK]

10894 The argument *cmd* is F\_SETLK or F\_SETLKW and satisfying the lock or unlock  
 10895 request would result in the number of locked regions in the system exceeding  
 a system-imposed limit.

10896 [EOVERFLOW] One of the values to be returned cannot be represented correctly.

10897 [EOVERFLOW]

10898 The *cmd* argument is F\_GETLK, F\_SETLK, or F\_SETLKW and the smallest or,  
 10899 if *l\_len* is non-zero, the largest offset of any byte in the requested segment  
 cannot be represented correctly in an object of type **off\_t**.

10900 The *fcntl()* function may fail if:

10901 [EDEADLK]

10902 The *cmd* argument is F\_SETLKW, the lock is blocked by some lock from  
 10903 another process and putting the calling process to sleep, waiting for that lock  
 to become free would cause a deadlock.

10904 **EXAMPLES**

10905 None.

10906 **APPLICATION USAGE**

10907 None.

10908 **RATIONALE**

10909 The ellipsis in the SYNOPSIS is the syntax specified by the ISO C standard for a variable number  
 10910 of arguments. It is used because System V uses pointers for the implementation of file locking  
 10911 functions.

10912 The *arg* values to F\_GETFD, F\_SETFD, F\_GETFL, and F\_SETFL all represent flag values to allow  
 10913 for future growth. Applications using these functions should do a read-modify-write operation  
 10914 on them, rather than assuming that only the values defined by this volume of  
 10915 IEEE Std. 1003.1-200x are valid. It is a common error to forget this, particularly in the case of  
 10916 F\_SETFD.

10917 This volume of IEEE Std. 1003.1-200x permits concurrent read and write access to file data using  
10918 the *fcntl()* function; this is a change from the 1984 /usr/group standard and early proposals.  
10919 Without concurrency controls, this feature may not be fully utilized without occasional loss of  
10920 data.

10921 Data losses occur in several ways. One case occurs when several processes try to update the  
10922 same record, without sequencing controls; several updates may occur in parallel and the last  
10923 writer “wins”. Another case is a bit-tree or other internal list-based database that is undergoing  
10924 reorganization. Without exclusive use to the tree segment by the updating process, other reading  
10925 processes chance getting lost in the database when the index blocks are split, condensed,  
10926 inserted, or deleted. While *fcntl()* is useful for many applications, it is not intended to be overly  
10927 general and does not handle the bit-tree example well.

10928 This facility is only required for regular files because it is not appropriate for many devices such  
10929 as terminals and network connections.

10930 Since *fcntl()* works with “any file descriptor associated with that file, however it is obtained”,  
10931 the file descriptor may have been inherited through a *fork()* or *exec* operation and thus may  
10932 affect a file that another process also has open.

10933 The use of the open file description to identify what to lock requires extra calls and presents  
10934 problems if several processes are sharing an open file description, but there are too many  
10935 implementations of the existing mechanism for this volume of IEEE Std. 1003.1-200x to use  
10936 different specifications.

10937 Another consequence of this model is that closing any file descriptor for a given file (whether or  
10938 not it is the same open file description that created the lock) causes the locks on that file to be  
10939 relinquished for that process. Equivalently, any close for any file/process pair relinquishes the  
10940 locks owned on that file for that process. But note that while an open file description may be  
10941 shared through *fork()*, locks are not inherited through *fork()*. Yet locks may be inherited through  
10942 one of the *exec* functions.

10943 The identification of a machine in a network environment is outside of the scope of this volume  
10944 of IEEE Std. 1003.1-200x. Thus, an *\_l\_sysid* member, such as found in System V, is not included in  
10945 the locking structure.

10946 Before successful return from an *F\_SETLK* or *F\_SETLKW* request, the previous lock type for  
10947 each byte in the specified region shall be replaced by the new lock type. This can result in a  
10948 previously locked region being split into smaller regions. If this would cause the number of  
10949 regions being held by all processes in the system to exceed a system-imposed limit, the *fcntl()*  
10950 function shall return  $-1$  with *errno* set to [ENOLCK].

10951 Mandatory locking was a major feature of the 1984 /usr/group standard.

10952 For advisory file record locking to be effective, all processes that have access to a file must  
10953 cooperate and use the advisory mechanism before doing I/O on the file. Enforcement-mode  
10954 record locking is important when it cannot be assumed that all processes are cooperating. For  
10955 example, if one user uses an editor to update a file at the same time that a second user executes  
10956 another process that updates the same file and if only one of the two processes is using advisory  
10957 locking, the processes are not cooperating. Enforcement-mode record locking would protect  
10958 against accidental collisions.

10959 Secondly, advisory record locking requires a process using locking to bracket each I/O operation  
10960 with lock (or test) and unlock operations. With enforcement-mode file and record locking, a  
10961 process can lock the file once and unlock when all I/O operations have been completed.  
10962 Enforcement-mode record locking provides a base that can be enhanced; for example, with  
10963 sharable locks. That is, the mechanism could be enhanced to allow a process to lock a file so

- 10964 other processes could read it, but none of them could write it.
- 10965 Mandatory locks were omitted for several reasons:
- 10966 1. Mandatory lock setting was done by multiplexing the set-group-ID bit in most  
10967 implementations; this was confusing, at best.
  - 10968 2. The relationship to file truncation as supported in 4.2 BSD was not well specified.
  - 10969 3. Any publicly readable file could be locked by anyone. Many historical implementations  
10970 keep the password database in a publicly readable file. A malicious user could thus  
10971 prohibit logins. Another possibility would be to hold open a long-distance telephone line.
  - 10972 4. Some demand-paged historical implementations offer memory mapped files, and  
10973 enforcement cannot be done on that type of file.
- 10974 Since sleeping on a region is interrupted with any signal, *alarm()* may be used to provide a  
10975 timeout facility in applications requiring it. This is useful in deadlock detection. Because  
10976 implementation of full deadlock detection is not always feasible, the [EDEADLK] error was  
10977 made optional.
- 10978 **FUTURE DIRECTIONS**
- 10979 None.
- 10980 **SEE ALSO**
- 10981 *close()*, *exec*, *open()*, *sigaction()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<fcntl.h>**,  
10982 **<signal.h>**, **<unistd.h>**
- 10983 **CHANGE HISTORY**
- 10984 First released in Issue 1. Derived from Issue 1 of the SVID.
- 10985 **Issue 4**
- 10986 The **<sys/types.h>** and **<unistd.h>** headers are now marked as optional (OH); these headers do  
10987 not need to be included on XSI-conformant systems.
- 10988 In the DESCRIPTION, sentences describing behavior when *l\_len* is negative are marked as an  
10989 extension, and the description of locks is corrected to make it a requirement on the application.
- 10990 The following change is incorporated for alignment with the ISO POSIX-1 standard:
- 10991 • In the DESCRIPTION, the meaning of a successful F\_SETLK or F\_SETLKW request is  
10992 clarified, after a POSIX Request for Interpretation.
- 10993 **Issue 5**
- 10994 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
10995 Threads Extension.
- 10996 Large File Summit extensions are added.
- 10997 **Issue 6**
- 10998 In the SYNOPSIS, the inclusion of **<sys/types.h>** is no longer required.
- 10999 The following new requirements on POSIX implementations derive from alignment with the  
11000 Single UNIX Specification:
- 11001 • The requirement to include **<sys/types.h>** has been removed. Although **<sys/types.h>** was  
11002 required for conforming implementations of previous POSIX specifications, it was not  
11003 required for UNIX applications.
  - 11004 • In the DESCRIPTION, sentences describing behavior when *l\_len* is negative are now  
11005 mandated, and the description of unlock (F\_UNLOCK) when *l\_len* is non-negative is  
11006 mandated.

- 11007           • In the ERRORS section, the [EINVAL] error condition has the case mandated when the *cmd* is  
11008           invalid, and two [Eoverflow] error conditions are added.
- 11009           The F\_GETOWN and F\_SETOWN values are added for sockets.
- 11010           The following changes were made to align with the IEEE P1003.1a draft standard:
- 11011           • Clarification is added that the extent of the bytes locked is determined prior to the blocking  
11012           action.
- 11013           The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that  
11014           *fcntl()* results are unspecified for typed memory objects.
- 11015           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

11016 **NAME**11017 `fcvt` — convert a floating-point number to a string (**LEGACY**)11018 **SYNOPSIS**11019 XSI `#include <stdlib.h>`11020 `char *fcvt(double value, int ndigit, int *restrict decpt,`  
11021 `int *restrict sign);`

11022

11023 **DESCRIPTION**11024 Refer to `ecvt()`.

11025 **NAME**11026 fdatasync — synchronize the data of a file (**REALTIME**)11027 **SYNOPSIS**

11028 SIO #include &lt;unistd.h&gt;

11029 int fdatasync(int *fil*des);

11030

11031 **DESCRIPTION**11032 The *fdatasync()* function shall force all currently queued I/O operations associated with the file  
11033 indicated by file descriptor *fil*des to the synchronized I/O completion state.11034 The functionality is as described for *fsync()* (with the symbol `_POSIX_SYNCHRONIZED_IO`  
11035 defined), with the exception that all I/O operations shall be completed as defined for  
11036 synchronized I/O data integrity completion.11037 **RETURN VALUE**11038 If successful, the *fdatasync()* function shall return the value 0; otherwise, the function shall return  
11039 the value `-1` and set *errno* to indicate the error. If the *fdatasync()* function fails, outstanding I/O  
11040 operations are not guaranteed to have been completed.11041 **ERRORS**11042 The *fdatasync()* function shall fail if:11043 [EBADF] The *fil*des argument is not a valid file descriptor open for writing.

11044 [EINVAL] This implementation does not support synchronized I/O for this file.

11045 In the event that any of the queued I/O operations fail, *fdatasync()* shall return the error  
11046 conditions defined for *read()* and *write()*.11047 **EXAMPLES**

11048 None.

11049 **APPLICATION USAGE**

11050 None.

11051 **RATIONALE**

11052 None.

11053 **FUTURE DIRECTIONS**

11054 None.

11055 **SEE ALSO**11056 *ai*o\_ *fsync()*, *fcntl()*, *fsync()*, *open()*, *read()*, *write()*, the Base Definitions volume of  
11057 IEEE Std. 1003.1-200x, <unistd.h>11058 **CHANGE HISTORY**

11059 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

11060 **Issue 6**11061 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
11062 implementation does not support the Synchronized Input and Output option.11063 The *fdatasync()* function is marked as part of the Synchronized Input and Output option.

11064 **NAME**11065 fdetach — detach a name from a STREAMS-based file descriptor (**STREAMS**)11066 **SYNOPSIS**

11067 XSR #include &lt;stropts.h&gt;

11068 int fdetach(const char \*path);

11069

11070 **DESCRIPTION**

11071 The *fdetach()* function detaches a STREAMS-based file from the file to which it was attached by a  
 11072 previous call to *fattach()*. The *path* argument points to the path name of the attached STREAMS  
 11073 file. The application shall ensure that the process has appropriate privileges or be the owner of  
 11074 the file. A successful call to *fdetach()* shall cause all path names that named the attached  
 11075 STREAMS file to again name the file to which the STREAMS file was attached. All subsequent  
 11076 operations on *path* shall operate on the underlying file and not on the STREAMS file.

11077 All open file descriptions established while the STREAMS file was attached to the file referenced  
 11078 by *path* shall still refer to the STREAMS file after the *fdetach()* has taken effect.

11079 If there are no open file descriptors or other references to the STREAMS file, then a successful  
 11080 call to *fdetach()* shall have the same effect as performing the last *close()* on the attached file.

11081 **RETURN VALUE**

11082 Upon successful completion, *fdetach()* shall return 0; otherwise, it shall return -1 and set *errno* to  
 11083 indicate the error.

11084 **ERRORS**11085 The *fdetach()* function shall fail if:

- |       |                |                                                                                              |  |
|-------|----------------|----------------------------------------------------------------------------------------------|--|
| 11086 | [EACCES]       | Search permission is denied on a component of the path prefix.                               |  |
| 11087 | [EINVAL]       | The <i>path</i> argument names a file that is not currently attached.                        |  |
| 11088 | [ELOOP]        | A loop exists in symbolic links encountered during resolution of the <i>path</i>             |  |
| 11089 |                | argument.                                                                                    |  |
| 11090 | [ENAMETOOLONG] |                                                                                              |  |
| 11091 |                | The size of a path name exceeds {PATH_MAX} or a path name component is                       |  |
| 11092 |                | longer than {NAME_MAX}.                                                                      |  |
| 11093 | [ENOENT]       | A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string. |  |
| 11094 | [ENOTDIR]      | A component of the path prefix is not a directory.                                           |  |
| 11095 | [EPERM]        | The effective user ID is not the owner of <i>path</i> and the process does not have          |  |
| 11096 |                | appropriate privileges.                                                                      |  |

11097 The *fdetach()* function may fail if:

- |       |                |                                                                         |  |
|-------|----------------|-------------------------------------------------------------------------|--|
| 11098 | [ELOOP]        | More than {SYMLOOP_MAX} symbolic links were encountered during          |  |
| 11099 |                | resolution of the <i>path</i> argument.                                 |  |
| 11100 | [ENAMETOOLONG] |                                                                         |  |
| 11101 |                | Path name resolution of a symbolic link produced an intermediate result |  |
| 11102 |                | whose length exceeds {PATH_MAX}.                                        |  |

11103 **EXAMPLES**11104 **Detaching a File**

11105 The following example detaches the STREAMS-based file **/tmp/named-STREAM** from the file to  
11106 which it was attached by a previous, successful call to *fattach()*. Subsequent calls to open this  
11107 file refer to the underlying file, not to the STREAMS file.

```
11108 #include <stropts.h>
11109 ...
11110 char *filename = "/tmp/named-STREAM";
11111 int ret;
11112 ret = fdetach(filename);
```

11113 **APPLICATION USAGE**

11114 None.

11115 **RATIONALE**

11116 None.

11117 **FUTURE DIRECTIONS**

11118 None.

11119 **SEE ALSO**

11120 *fattach()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stropts.h>

11121 **CHANGE HISTORY**

11122 First released in Issue 4, Version 2.

11123 **Issue 5**

11124 Moved from X/OPEN UNIX extension to BASE.

11125 **Issue 6**

11126 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

11127 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
11128 [ELOOP] error condition is added.

11129 **NAME**

11130 fdim, fdimf, fdiml — compute positive difference between two floating-point numbers

11131 **SYNOPSIS**

11132 #include <math.h>

11133 double fdim(double x, double y);

11134 float fdimf(float x, float y);

11135 long double fdiml(long double x, long double y);

11136 **DESCRIPTION**

11137 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
11138 conflict between the requirements described here and the ISO C standard is unintentional. This  
11139 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11140 These functions shall determine the positive difference between their arguments. If *x* is greater  
11141 than *y*, *x*−*y* is returned. If *x* is less than or equal to *y*, +0 is returned.

11142 An application wishing to check for error situations should set *errno* to 0 before calling these  
11143 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

11144 **RETURN VALUE**

11145 Upon successful completion, these functions shall return the positive difference value.

11146 If *x* or *y* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

11147 If the magnitude of the result is too large or too small, the numeric result is unspecified and *errno*  
11148 may be set to [ERANGE].

11149 **ERRORS**

11150 These functions may fail if:

11151 [EDOM] The value of *x* or *y* is NaN.

11152 [ERANGE] The magnitude of the result is too large or too small.

11153 **EXAMPLES**

11154 None.

11155 **APPLICATION USAGE**

11156 None.

11157 **RATIONALE**

11158 None.

11159 **FUTURE DIRECTIONS**

11160 None.

11161 **SEE ALSO**

11162 *fmax()*, *fmin()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

11163 **CHANGE HISTORY**

11164 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11165 **NAME**

11166 fdopen — associate a stream with a file descriptor

11167 **SYNOPSIS**

11168 #include &lt;stdio.h&gt;

11169 FILE \*fdopen(int *fil-des*, const char \**mode*);11170 **DESCRIPTION**11171 The *fdopen()* function shall associate a stream with a file descriptor.11172 The *mode* argument is a character string having one of the following values:11173 *r* or *rb* Open a file for reading.11174 *w* or *wb* Open a file for writing.11175 *a* or *ab* Open a file for writing at end of file.11176 *r+* or *rb+* or *r+b* Open a file for update (reading and writing).11177 *w+* or *wb+* or *w+b* Open a file for update (reading and writing).11178 *a+* or *ab+* or *a+b* Open a file for update (reading and writing) at end of file.11179 The meaning of these flags is exactly as specified in *fopen()*, except that modes beginning with *w*  
11180 do not cause truncation of the file.11181 Additional values for the *mode* argument may be supported by an implementation.11182 The application shall ensure that the mode of the stream as expressed by the *type* argument is  
11183 allowed by the file access mode of the open file description to which *fil-des* refers. The file  
11184 position indicator associated with the new stream is set to the position indicated by the file  
11185 offset associated with the file descriptor.11186 The error and end-of-file indicators for the stream shall be cleared. The *fdopen()* function may  
11187 cause the *st\_atime* field of the underlying file to be marked for update.11188 SHM If *fil-des* refers to a shared memory object, the result of the *fdopen()* function is unspecified.11189 TYM If *fil-des* refers to a typed memory object, the result of the *fdopen()* function is unspecified.11190 The *fdopen()* function shall preserve the offset maximum previously set for the open file  
11191 description corresponding to *fil-des*.11192 **RETURN VALUE**11193 Upon successful completion, *fdopen()* shall return a pointer to a stream; otherwise, a null pointer  
11194 shall be returned and *errno* set to indicate the error.11195 **ERRORS**11196 The *fdopen()* function may fail if:11197 [EBADF] The *fil-des* argument is not a valid file descriptor. |11198 [EINVAL] The *mode* argument is not a valid mode. |

11199 [EMFILE] {FOPEN\_MAX} streams are currently open in the calling process. |

11200 [EMFILE] {STREAM\_MAX} streams are currently open in the calling process. |

11201 [ENOMEM] Insufficient space to allocate a buffer. |

11202 **EXAMPLES**

11203 None.

11204 **APPLICATION USAGE**

11205 File descriptors are obtained from calls like *open()*, *dup()*, *creat()*, or *pipe()*, which open files but  
 11206 do not return streams.

11207 **RATIONALE**

11208 The file descriptor may have been obtained from *open()*, *creat()*, *pipe()*, *dup()*, or *fcntl()*;  
 11209 inherited through *fork()* or *exec*; or perhaps obtained by implementation-defined means, such as  
 11210 the 4.3 BSD *socket()* call.

11211 The meanings of the *type* arguments of *fdopen()* and *fopen()* differ. With *fdopen()*, open for write  
 11212 (*w* or *w+*) does not truncate, and append (*a* or *a+*) cannot create for writing. There is no need for *b*  
 11213 in the format due to the equivalence of binary and text files in this volume of  
 11214 IEEE Std. 1003.1-200x. Although not explicitly required by this volume of IEEE Std. 1003.1-200x,  
 11215 a good implementation of append (*a*) mode would cause the O\_APPEND flag to be set.

11216 **FUTURE DIRECTIONS**

11217 None.

11218 **SEE ALSO**

11219 *fclose()*, *fopen()*, *open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdio.h**>, Section  
 11220 2.5.1 (on page 535)

11221 **CHANGE HISTORY**

11222 First released in Issue 1. Derived from Issue 1 of the SVID.

11223 **Issue 4**

11224 In the DESCRIPTION, the use and settings of the *mode* argument are changed to include binary  
 11225 streams and are marked as extensions.

11226 All errors identified in the ERRORS section are marked as extensions, and the [EMFILE] error is  
 11227 added.

11228 The APPLICATION USAGE section is added.

11229 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 11230 • The type of argument *mode* is changed from **char\*** to **const char\***.

11231 **Issue 5**

11232 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.

11233 Large File Summit extensions are added.

11234 **Issue 6**

11235 The following new requirements on POSIX implementations derive from alignment with the  
 11236 Single UNIX Specification:

- 11237 • In the DESCRIPTION, the use and setting of the *mode* argument are changed to include  
 11238 binary streams.

- 11239 • In the DESCRIPTION, text is added for large file support to indicate setting of the offset  
 11240 maximum in the open file description.

- 11241 • All errors identified in the ERRORS section are added.

- 11242 • In the DESCRIPTION, text is added that the *fdopen()* function may cause *st\_atime* to be  
 11243 updated.

- 11244 The following changes were made to align with the IEEE P1003.1a draft standard:
- 11245 • Clarification is added that it is the responsibility of the application to ensure that the mode is  
11246 compatible with the open file descriptor.
- 11247 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that |  
11248 *fdopen()* results are unspecified for typed memory objects.

11249 **NAME**

11250        `feclearexcept` — clear floating-point exception

11251 **SYNOPSIS**

11252        `#include <fenv.h>`

11253        `void feclearexcept(int excepts);`

11254 **DESCRIPTION**

11255 *cx*        The functionality described on this reference page is aligned with the ISO C standard. Any  
11256        conflict between the requirements described here and the ISO C standard is unintentional. This  
11257        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11258        The `feclearexcept()` function shall clear the supported floating-point exceptions represented by  
11259        *excepts*.

11260 **RETURN VALUE**

11261        None.

11262 **ERRORS**

11263        No errors are defined.

11264 **EXAMPLES**

11265        None.

11266 **APPLICATION USAGE**

11267        None.

11268 **RATIONALE**

11269        None.

11270 **FUTURE DIRECTIONS**

11271        None.

11272 **SEE ALSO**

11273        `fegetexceptflag()`, `feraiseexcept()`, `fesetexceptflag()`, `fetestexcept()`, the Base Definitions volume of  
11274        IEEE Std. 1003.1-200x, `<fenv.h>`

11275 **CHANGE HISTORY**

11276        First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11277 **NAME**

11278 fegetenv, fesetenv — get and set current floating-point environment

11279 **SYNOPSIS**

11280 #include <fenv.h>

11281 void fegetenv(fenv\_t \*envp);

11282 void fesetenv(const fenv\_t \*envp);

11283 **DESCRIPTION**

11284 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
11285 conflict between the requirements described here and the ISO C standard is unintentional. This  
11286 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11287 The *fegetenv()* function shall store the current floating-point environment in the object pointed to  
11288 by *envp*.

11289 The *fesetenv()* function shall establish the floating-point environment represented by the object  
11290 pointed to by *envp*. The argument *envp* shall point to an object set by a call to *fegetenv()* or  
11291 *fehldexcept()*, or equal a floating-point environment macro. The *fesetenv()* function does not  
11292 raise floating-point exceptions, but only installs the state of the floating-point status flags  
11293 represented through its argument.

11294 **RETURN VALUE**

11295 None.

11296 **ERRORS**

11297 No errors are defined.

11298 **EXAMPLES**

11299 None.

11300 **APPLICATION USAGE**

11301 None.

11302 **RATIONALE**

11303 None.

11304 **FUTURE DIRECTIONS**

11305 None.

11306 **SEE ALSO**

11307 *fehldexcept()*, *feupdateenv()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <fenv.h>

11308 **CHANGE HISTORY**

11309 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11310 **NAME**

11311 fegetexceptflag, fesetexceptflag — get and set floating-point status flags

11312 **SYNOPSIS**

11313 #include <fenv.h>

11314 void fegetexceptflag(fexcept\_t \*flagp, int excepts);

11315 void fesetexceptflag(const fexcept\_t \*flagp, int excepts);

11316 **DESCRIPTION**

11317 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
11318 conflict between the requirements described here and the ISO C standard is unintentional. This  
11319 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11320 The *fegetexceptflag()* function shall store an implementation-defined representation of the states  
11321 of the floating-point status flags indicated by the argument *excepts* in the object pointed to by the  
11322 argument *flagp*.

11323 The *fesetexceptflag()* function shall set the floating-point status flags indicated by the argument  
11324 *excepts* to the states stored in the object pointed to by *flagp*. The value pointed to by *flagp* shall  
11325 have been set by a previous call to *fegetexceptflag()* whose second argument represented at least  
11326 those floating-point exceptions represented by the argument *excepts*. This function does not  
11327 raise floating-point exceptions, but only sets the state of the flags.

11328 **RETURN VALUE**

11329 None.

11330 **ERRORS**

11331 No errors are defined.

11332 **EXAMPLES**

11333 None.

11334 **APPLICATION USAGE**

11335 None.

11336 **RATIONALE**

11337 None.

11338 **FUTURE DIRECTIONS**

11339 None.

11340 **SEE ALSO**

11341 *feclearexcept()*, *feraiseexcept()*, *fetestexcept()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
11342 <fenv.h>

11343 **CHANGE HISTORY**

11344 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11345 **NAME**

11346 fegetround, fesetround — get and set current rounding direction

11347 **SYNOPSIS**

```
11348 #include <fenv.h>
11349 int fegetround(void);
11350 int fesetround(int round);
```

11351 **DESCRIPTION**

11352 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 11353 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11354 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11355 The *fegetround()* function shall get the current rounding direction.

11356 The *fesetround()* function shall establish the rounding direction represented by its argument  
 11357 *round*. If the argument is not equal to the value of a rounding direction macro, the rounding  
 11358 direction is not changed.

11359 **RETURN VALUE**

11360 The *fegetround()* function shall return the value of the rounding direction macro representing the  
 11361 current rounding direction or a negative value if there is no such rounding direction macro or  
 11362 the current rounding direction is not determinable.

11363 The *fesetround()* function shall return a zero value if and only if the argument is equal to a  
 11364 rounding direction macro (that is, if and only if the requested rounding direction was  
 11365 established).

11366 **ERRORS**

11367 No errors are defined.

11368 **EXAMPLES**

11369 The following example saves, sets, and restores the rounding direction, reporting an error and  
 11370 aborting if setting the rounding direction fails:

```
11371 #include <fenv.h>
11372 #include <assert.h>
11373 void f(int round_dir)
11374 {
11375 #pragma STDC FENV_ACCESS ON
11376 int save_round;
11377 int setround_ok;
11378 save_round = fegetround();
11379 setround_ok = fesetround(round_dir);
11380 assert(setround_ok == 0);
11381 /* ... */
11382 fesetround(save_round);
11383 /* ... */
11384 }
```

11385 **APPLICATION USAGE**

11386 None.

11387 **RATIONALE**

11388 None.

11389 **FUTURE DIRECTIONS**

11390           None.

11391 **SEE ALSO**

11392           The Base Definitions volume of IEEE Std. 1003.1-200x, <**fenv.h**>

11393 **CHANGE HISTORY**

11394           First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard.

11395 **NAME**

11396 feholdexcept — save current floating-point environment

11397 **SYNOPSIS**

11398 #include <fenv.h>

11399 int feholdexcept(fenv\_t \*envp);

11400 **DESCRIPTION**

11401 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
11402 conflict between the requirements described here and the ISO C standard is unintentional. This  
11403 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11404 The *feholdexcept()* function shall save the current floating-point environment in the object  
11405 pointed to by *envp*, clear the floating-point status flags, and then install a non-stop (continue on  
11406 floating-point exceptions) mode, if available, for all floating-point exceptions.

11407 **RETURN VALUE**

11408 The *feholdexcept()* function shall return zero if and only if non-stop floating-point exception  
11409 handling was successfully installed.

11410 **ERRORS**

11411 No errors are defined.

11412 **EXAMPLES**

11413 None.

11414 **APPLICATION USAGE**

11415 None.

11416 **RATIONALE**

11417 The *feholdexcept()* function should be effective on typical IEC 60559:1989 standard  
11418 implementations which have the default non-stop mode and at least one other mode for trap  
11419 handling or aborting. If the implementation provides only the non-stop mode, then installing the  
11420 non-stop mode is trivial.

11421 **FUTURE DIRECTIONS**

11422 None.

11423 **SEE ALSO**

11424 *fegetenv()*, *fesetenv()*, *feupdateenv()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
11425 <fenv.h>

11426 **CHANGE HISTORY**

11427 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11428 **NAME**

11429 feof — test end-of-file indicator on a stream

11430 **SYNOPSIS**

11431 #include &lt;stdio.h&gt;

11432 int feof(FILE \*stream);

11433 **DESCRIPTION**

11434 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
11435 conflict between the requirements described here and the ISO C standard is unintentional. This  
11436 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11437 The *feof()* function shall test the end-of-file indicator for the stream pointed to by *stream*.11438 **RETURN VALUE**11439 The *feof()* function shall return non-zero if and only if the end-of-file indicator is set for *stream*.11440 **ERRORS**

11441 No errors are defined.

11442 **EXAMPLES**

11443 None.

11444 **APPLICATION USAGE**

11445 None.

11446 **RATIONALE**

11447 None.

11448 **FUTURE DIRECTIONS**

11449 None.

11450 **SEE ALSO**11451 *clearerr()*, *ferror()*, *fopen()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>11452 **CHANGE HISTORY**

11453 First released in Issue 1. Derived from Issue 1 of the SVID.

11454 **Issue 4**

11455 The **ERRORS** section is rewritten, such that no error return values are now defined for this  
11456 function.

11457 **NAME**

11458        `feraiseexcept` — raise floating-point exception

11459 **SYNOPSIS**

11460        `#include <fenv.h>`

11461        `void feraiseexcept(int excepts);`

11462 **DESCRIPTION**

11463 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
11464        conflict between the requirements described here and the ISO C standard is unintentional. This  
11465        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11466        The *feraiseexcept()* function shall raise the supported floating-point exceptions represented by  
11467        the argument *excepts*. The order in which these floating-point exceptions are raised is  
11468        unspecified. Whether the *feraiseexcept()* function additionally raises the inexact floating-point  
11469        exception whenever it raises the overflow or underflow floating-point exception is  
11470        implementation-defined.

11471 **RETURN VALUE**

11472        None.

11473 **ERRORS**

11474        No errors are defined.

11475 **EXAMPLES**

11476        None.

11477 **APPLICATION USAGE**

11478        The effect is intended to be similar to that of floating-point exceptions raised by arithmetic  
11479        operations. Hence, enabled traps for floating-point exceptions raised by this function are taken.

11480 **RATIONALE**

11481        Raising overflow or underflow is allowed to also raise inexact because on some architectures the  
11482        only practical way to raise an exception is to execute an instruction that has the exception as a  
11483        side effect. The function is not restricted to accept only valid coincident expressions for atomic  
11484        operations, so the function can be used to raise exceptions accrued over several operations.

11485 **FUTURE DIRECTIONS**

11486        None.

11487 **SEE ALSO**

11488        *feclearexcept()*, *fegetexceptflag()*, *fesetexceptflag()*, *fetestexcept()*, the Base Definitions volume of  
11489        IEEE Std. 1003.1-200x, `<fenv.h>`

11490 **CHANGE HISTORY**

11491        First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11492 **NAME**

11493       ferror — test error indicator on a stream

11494 **SYNOPSIS**

11495       #include &lt;stdio.h&gt;

11496       int ferror(FILE \**stream*);11497 **DESCRIPTION**

11498 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
11499       conflict between the requirements described here and the ISO C standard is unintentional. This  
11500       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11501       The *ferror()* function shall test the error indicator for the stream pointed to by *stream*.11502 **RETURN VALUE**11503       The *ferror()* function shall return non-zero if and only if the error indicator is set for *stream*.11504 **ERRORS**

11505       No errors are defined.

11506 **EXAMPLES**

11507       None.

11508 **APPLICATION USAGE**

11509       None.

11510 **RATIONALE**

11511       None.

11512 **FUTURE DIRECTIONS**

11513       None.

11514 **SEE ALSO**11515       *clearerr()*, *feof()*, *fopen()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>11516 **CHANGE HISTORY**

11517       First released in Issue 1. Derived from Issue 1 of the SVID.

11518 **Issue 4**

11519       The **ERRORS** section is rewritten, such that no error return values are now defined for this  
11520       function.

11521 **NAME**

11522       fesetenv — set current floating-point environment

11523 **SYNOPSIS**

11524       #include &lt;fenv.h&gt;

11525       void fesetenv(const fenv\_t \*envp);

11526 **DESCRIPTION**11527       Refer to *fegetenv()*.

11528 **NAME**

11529       fesetexceptflag — set floating-point status flags

11530 **SYNOPSIS**

11531       #include <fenv.h>

11532       void fesetexceptflag(const fexcept\_t \*flagp, int excepts);

11533 **DESCRIPTION**

11534       Refer to *fegetexceptflag()*.

11535 **NAME**

11536       fesetround — set current rounding direction

11537 **SYNOPSIS**

11538       #include &lt;fenv.h&gt;

11539       int fesetround(int *round*);11540 **DESCRIPTION**11541       Refer to *fegetround()*.

11542 **NAME**

11543       fetestexcept — test floating-point exception flags

11544 **SYNOPSIS**

11545       #include &lt;fenv.h&gt;

11546       int fetestexcept(int *excepts*);11547 **DESCRIPTION**

11548 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
11549 conflict between the requirements described here and the ISO C standard is unintentional. This  
11550 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11551       The *fetestexcept()* function shall determine which of a specified subset of the floating-point  
11552 exception flags are currently set. The *excepts* argument specifies the floating-point status flags to  
11553 be queried.

11554 **RETURN VALUE**

11555       The *fetestexcept()* function shall return the value of the bitwise-inclusive OR of the floating-point  
11556 exception macros corresponding to the currently set floating-point exceptions included in  
11557 *excepts*.

11558 **ERRORS**

11559       No errors are defined.

11560 **EXAMPLES**

11561       The following example calls function *f()* if an invalid exception is set, and then function *g()* if an  
11562 overflow exception is set:

```
11563 #include <fenv.h>
11564 /* ... */
11565 {
11566 #pragma STDC FENV_ACCESS ON
11567 int set_excepts;
11568 feclearexcept(FE_INVALID | FE_OVERFLOW);
11569 // maybe raise exceptions
11570 set_excepts = fetestexcept(FE_INVALID | FE_OVERFLOW);
11571 if (set_excepts & FE_INVALID) f();
11572 if (set_excepts & FE_OVERFLOW) g();
11573 /* ... */
11574 }
```

11575 **APPLICATION USAGE**

11576       None.

11577 **RATIONALE**

11578       None.

11579 **FUTURE DIRECTIONS**

11580       None.

11581 **SEE ALSO**

11582       *feclearexcept()*, *fegetexceptflag()*, *feraiseexcept()*, the Base Definitions volume of  
11583 IEEE Std. 1003.1-200x, <fenv.h>

11584 **CHANGE HISTORY**

11585 First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard.

11586 **NAME**

11587 feupdateenv — update floating-point environment

11588 **SYNOPSIS**

11589 #include &lt;fenv.h&gt;

11590 void feupdateenv(const fenv\_t \*envp);

11591 **DESCRIPTION**

11592 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
11593 conflict between the requirements described here and the ISO C standard is unintentional. This  
11594 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11595 The *feupdateenv()* function shall save the currently raised floating-point exceptions in its  
11596 automatic storage, install the floating-point environment represented by the object pointed to by  
11597 *envp*, and then raise the saved floating-point exceptions. The argument *envp* shall point to an  
11598 object set by a call to *feholdexcept()* or *fegetenv()*, or equal a floating-point environment macro.

11599 **RETURN VALUE**

11600 None.

11601 **ERRORS**

11602 No errors are defined.

11603 **EXAMPLES**

11604 The following example shows sample code to hide spurious underflow floating-point  
11605 exceptions:

```
11606 #include <fenv.h>
11607 double f(double x)
11608 {
11609 #pragma STDC FENV_ACCESS ON
11610 double result;
11611 fenv_t save_env;
11612 feholdexcept(&save_env);
11613 // compute result
11614 if (/* test spurious underflow */)
11615 feclearexcept(FE_UNDERFLOW);
11616 feupdateenv(&save_env);
11617 return result;
11618 }
```

11619 **APPLICATION USAGE**

11620 None.

11621 **RATIONALE**

11622 None.

11623 **FUTURE DIRECTIONS**

11624 None.

11625 **SEE ALSO**11626 *fegetenv()*, *feholdexcept()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <fenv.h>11627 **CHANGE HISTORY**

11628 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

11629 **NAME**

11630 fflush — flush a stream

11631 **SYNOPSIS**

11632 #include &lt;stdio.h&gt;

11633 int fflush(FILE \**stream*);11634 **DESCRIPTION**

11635 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 11636 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11637 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11638 If *stream* points to an output stream or an update stream in which the most recent operation was  
 11639 CX not input, *fflush()* causes any unwritten data for that stream to be written to the file, and the  
 11640 *st\_ctime* and *st\_mtime* fields of the underlying file are marked for update.

11641 If *stream* is a null pointer, *fflush()* shall perform this flushing action on all streams for which the  
 11642 behavior is defined above.

11643 **RETURN VALUE**

11644 Upon successful completion, *fflush()* shall return 0; otherwise, it shall set the error indicator for  
 11645 CX the stream, return EOF, and set *errno* to indicate the error.

11646 **ERRORS**11647 The *fflush()* function shall fail if:

11648 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor underlying *stream* and the  
 11649 process would be delayed in the write operation.

11650 CX [EBADF] The file descriptor underlying *stream* is not valid.

11651 CX [EFBIG] An attempt was made to write a file that exceeds the maximum file size.

11652 XSI [EFBIG] An attempt was made to write a file that exceeds the process' file size limit.

11653 CX [EFBIG] The file is a regular file and an attempt was made to write at or beyond the  
 11654 offset maximum associated with the corresponding stream.

11655 CX [EINTR] The *fflush()* function was interrupted by a signal.

11656 CX [EIO] The process is a member of a background process group attempting to write  
 11657 to its controlling terminal, TOSTOP is set, the process is neither ignoring nor  
 11658 blocking SIGTTOU, and the process group of the process is orphaned. This  
 11659 error may also be returned under implementation-defined conditions.

11660 CX [ENOSPC] There was no free space remaining on the device containing the file.

11661 CX [EPIPE] An attempt is made to write to a pipe or FIFO that is not open for reading by  
 11662 any process. A SIGPIPE signal shall also be sent to the thread.

11663 The *fflush()* function may fail if:

11664 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
 11665 capabilities of the device.

11666 **EXAMPLES**11667 **Sending Prompts to Standard Output**

11668 The following example uses *printf()* calls to print a series of prompts for information the user  
 11669 must enter from standard input. The *fflush()* calls force the output to standard output. The  
 11670 *fflush()* function is used because standard output is usually buffered and the prompt may not  
 11671 immediately be printed on the output or terminal. The *gets()* calls read strings from standard  
 11672 input and place the results in variables, for use later in the program

```
11673 #include <stdio.h>
11674 ...
11675 char user[100];
11676 char oldpasswd[100];
11677 char newpasswd[100];
11678 ...
11679 printf("User name: ");
11680 fflush(stdout);
11681 gets(user);

11682 printf("Old password: ");
11683 fflush(stdout);
11684 gets(oldpasswd);

11685 printf("New password: ");
11686 fflush(stdout);
11687 gets(newpasswd);
11688 ...
```

11689 **APPLICATION USAGE**

11690 None.

11691 **RATIONALE**

11692 Data buffered by the system may make determining the validity of the position of the current  
 11693 file descriptor impractical. Thus, enforcing the repositioning of the file descriptor after *fflush()*  
 11694 on streams open for *read()* is not mandated by IEEE Std. 1003.1-200x.

11695 **FUTURE DIRECTIONS**

11696 None.

11697 **SEE ALSO**

11698 *getrlimit()*, *ulimit()*, the Base Definitions volume of IEEE Std. 1003.1-200x, *<stdio.h>*

11699 **CHANGE HISTORY**

11700 First released in Issue 1. Derived from Issue 1 of the SVID.

11701 **Issue 4**

11702 The following change is incorporated for alignment with the ISO C standard:

- 11703 • The DESCRIPTION is changed to describe the behavior of *fflush()* if *stream* is a null pointer.

11704 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 11705 • The following two paragraphs are withdrawn from the DESCRIPTION (by POSIX as well as  
 11706 X/Open) because of the possibility of causing applications to malfunction, and the  
 11707 impossibility of implementing these mechanisms for pipes:

- 11708 If the stream is open for reading, any unread data buffered in the stream is discarded.
- 11709 For a stream open for reading, if the file is not already at EOF, and the file is one capable  
11710 of seeking, the file offset of the underlying open file description is adjusted so that the  
11711 next operation on the open file description deals with the byte after the last one read  
11712 from, or written to, the stream being flushed.
- 11713
  - The [EFBIG] error is marked to indicate the extensions.
- 11714 **Issue 5**
- 11715 Large File Summit extensions are added.
- 11716 **Issue 6**
- 11717 Extensions beyond the ISO C standard are now marked.
- 11718 The following new requirements on POSIX implementations derive from alignment with the  
11719 Single UNIX Specification:
- 11720
  - The [EFBIG] error is added as part of the large file support extensions.
- 11721
  - The [ENXIO] optional error condition is added.
- 11722 The RETURN VALUE section is updated to note that the error indicator shall be set for the  
11723 stream. This is for alignment with the ISO/IEC 9899:1999 standard. |

11724 **NAME**

11725 ffs — find first set bit

11726 **Notes to Reviewers**11727 *This section with side shading will not appear in the final copy. - Ed.*11728 This function or these functions are recommended to become mandatory parts of POSIX.1 in the  
11729 next draft.11730 **SYNOPSIS**11731 xSI `#include <strings.h>`11732 `int ffs(int i);`

11733

11734 **DESCRIPTION**11735 The `ffs()` function shall find the first bit set (beginning with the least significant bit) in *i*, and  
11736 return the index of that bit. Bits are numbered starting at one (the least significant bit).11737 **RETURN VALUE**11738 The `ffs()` function shall return the index of the first bit set. If *i* is 0, then `ffs()` shall return 0.11739 **ERRORS**

11740 No errors are defined.

11741 **EXAMPLES**

11742 None.

11743 **APPLICATION USAGE**

11744 None.

11745 **RATIONALE**

11746 None.

11747 **FUTURE DIRECTIONS**

11748 None.

11749 **SEE ALSO**11750 The Base Definitions volume of IEEE Std. 1003.1-200x, `<strings.h>`11751 **CHANGE HISTORY**

11752 First released in Issue 4, Version 2.

11753 **Issue 5**

11754 Moved from X/OPEN UNIX extension to BASE.

11755 **NAME**

11756 fgetc — get a byte from a stream

11757 **SYNOPSIS**

11758 #include &lt;stdio.h&gt;

11759 int fgetc(FILE \**stream*);11760 **DESCRIPTION**

11761 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 11762 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11763 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11764 If the end-of-file indicator for the input stream pointed to by *stream* is not set and a next  
 11765 character is present, the *fgetc()* function obtains the next byte (if present) as an **unsigned char**  
 11766 converted to an **int**, from the input stream pointed to by *stream*, and advances the associated file  
 11767 position indicator for the stream (if defined).

11768 CX The *fgetc()* function may mark the *st\_atime* field of the file associated with *stream* for update. The  
 11769 *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
 11770 *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns  
 11771 data not supplied by a prior call to *ungetc()* or *ungetwc()*.

11772 **RETURN VALUE**

11773 Upon successful completion, *fgetc()* shall return the next byte from the input stream pointed to  
 11774 by *stream*. If the end-of-file indicator for the stream is set, or if the stream is at end-of-file, the  
 11775 end-of-file indicator for the stream shall be set and *fgetc()* shall return EOF. If a read error occurs,  
 11776 CX the error indicator for the stream shall be set, *fgetc()* shall return EOF, and shall set *errno* to  
 11777 indicate the error.

11778 **ERRORS**11779 The *fgetc()* function shall fail if data needs to be read and:

11780 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor underlying *stream* and the  
 11781 process would be delayed in the *fgetc()* operation.

11782 CX [EBADF] The file descriptor underlying *stream* is not a valid file descriptor open for  
 11783 reading.

11784 CX [EINTR] The read operation was terminated due to the receipt of a signal, and no data  
 11785 was transferred.

11786 CX [EIO] A physical I/O error has occurred, or the process is in a background process  
 11787 group attempting to read from its controlling terminal, and either the process  
 11788 is ignoring or blocking the SIGTTIN signal or the process group is orphaned.  
 11789 This error may also be generated for implementation-defined reasons.

11790 CX [EOVERFLOW] The file is a regular file and an attempt was made to read at or beyond the  
 11791 offset maximum associated with the corresponding stream.

11792 The *fgetc()* function may fail if:

11793 CX [ENOMEM] Insufficient storage space is available.

11794 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
 11795 capabilities of the device.

11796 **EXAMPLES**

11797 None.

11798 **APPLICATION USAGE**

11799 If the integer value returned by *fgetc()* is stored into a variable of type **char** and then compared  
11800 against the integer constant EOF, the comparison may never succeed, because sign-extension of  
11801 a variable of type **char** on widening to integer is implementation-defined.

11802 The *ferror()* or *feof()* functions must be used to distinguish between an error condition and an  
11803 end-of-file condition.

11804 **RATIONALE**

11805 None.

11806 **FUTURE DIRECTIONS**

11807 None.

11808 **SEE ALSO**

11809 *feof()*, *ferror()*, *fopen()*, *getchar()*, *getc()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
11810 **<stdio.h>**

11811 **CHANGE HISTORY**

11812 First released in Issue 1. Derived from Issue 1 of the SVID.

11813 **Issue 4**

11814 In the DESCRIPTION:

- 11815 • The text is changed to make it clear that the function returns a byte value.
- 11816 • The list of functions that may cause the *st\_atime* field to be updated is revised.

11817 In the ERRORS section, text is added to indicate that error returns are only generated when data  
11818 needs to be read into the stream buffer.

11819 Also in the ERRORS section, in previous issues generation of the [EIO] error depended on  
11820 whether or not an implementation supported Job Control. This functionality is now defined as  
11821 mandatory.

11822 The [ENXIO] and [ENOMEM] errors are marked as extensions.

11823 In the APPLICATION USAGE section, text is added to indicate how an application can  
11824 distinguish between an error condition and an end-of-file condition.

11825 The description of [EINTR] is amended.

11826 **Issue 4, Version 2**

11827 In the ERRORS section, the description of [EIO] is updated to include the case where a physical  
11828 I/O error occurs.

11829 **Issue 5**

11830 Large File Summit extensions are added.

11831 **Issue 6**

11832 Extensions beyond the ISO C standard are now marked.

11833 The following new requirements on POSIX implementations derive from alignment with the  
11834 Single UNIX Specification:

- 11835 • The [EIO] and [EOVERFLOW] mandatory error conditions are added.
- 11836 • The [ENOMEM] and [ENXIO] optional error conditions are added.

- 11837 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:
- 11838 • The DESCRIPTION is updated to clarify the behavior when the end-of-file indicator for the  
11839 input stream is not set.
  - 11840 • The RETURN VALUE section is updated to note that the error indicator shall be set for the  
11841 stream.

11842 **NAME**

11843 fgetpos — get current file position information

11844 **SYNOPSIS**

11845 #include <stdio.h>

11846 int fgetpos(FILE \*restrict stream, fpos\_t \*restrict pos);

11847 **DESCRIPTION**

11848 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
11849 conflict between the requirements described here and the ISO C standard is unintentional. This  
11850 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11851 The *fgetpos()* function shall store the current value of the file position indicator for the stream  
11852 pointed to by *stream* in the object pointed to by *pos*. The value stored contains unspecified  
11853 information usable by *fsetpos()* for repositioning the stream to its position at the time of the call  
11854 to *fgetpos()*.

11855 **RETURN VALUE**

11856 Upon successful completion, *fgetpos()* shall return 0; otherwise, it shall return a non-zero value  
11857 and set *errno* to indicate the error.

11858 **ERRORS**

11859 The *fgetpos()* function shall fail if:

11860 CX [EOVERFLOW] The current value of the file position cannot be represented correctly in an  
11861 object of type **fpos\_t**.

11862 The *fgetpos()* function may fail if:

11863 CX [EBADF] The file descriptor underlying *stream* is not valid.

11864 CX [ESPIPE] The file descriptor underlying *stream* is associated with a pipe, FIFO, or socket.  
11865

11866 **EXAMPLES**

11867 None.

11868 **APPLICATION USAGE**

11869 None.

11870 **RATIONALE**

11871 None.

11872 **FUTURE DIRECTIONS**

11873 None.

11874 **SEE ALSO**

11875 *fopen()*, *ftell()*, *rewind()*, *ungetc()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
11876 <stdio.h>

11877 **CHANGE HISTORY**

11878 First released in Issue 4. Derived from the ISO C standard.

11879 **Issue 5**

11880 Large File Summit extensions are added.

11881 **Issue 6**

11882 Extensions beyond the ISO C standard are now marked.

11883 The following new requirements on POSIX implementations derive from alignment with the  
11884 Single UNIX Specification:

- 11885      • The [EIO] mandatory error condition is added.
- 11886      • The [EBADF] and [ESPIPE] optional error conditions are added.
- 11887      An additional [ESPIPE] error condition is added for sockets.
- 11888      The prototype for *fgetpos()* is changed for alignment with the ISO/IEC 9899:1999 standard.

11889 **NAME**

11890        fgets — get a string from a stream

11891 **SYNOPSIS**

11892        #include &lt;stdio.h&gt;

11893        char \*fgets(char \*restrict *s*, int *n*, FILE \*restrict *stream*);11894 **DESCRIPTION**

11895 cx        The functionality described on this reference page is aligned with the ISO C standard. Any  
 11896        conflict between the requirements described here and the ISO C standard is unintentional. This  
 11897        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11898        The *fgets()* function shall read bytes from *stream* into the array pointed to by *s*, until *n*−1 bytes  
 11899        are read, or a <newline> character is read and transferred to *s*, or an end-of-file condition is  
 11900        encountered. The string is then terminated with a null byte.

11901 cx        The *fgets()* function may mark the *st\_atime* field of the file associated with *stream* for update. The  
 11902        *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
 11903        *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns  
 11904        data not supplied by a prior call to *ungetc()* or *ungetwc()*.

11905 **RETURN VALUE**

11906        Upon successful completion, *fgets()* shall return *s*. If the stream is at end-of-file, the end-of-file  
 11907        indicator for the stream shall be set and *fgets()* shall return a null pointer. If a read error occurs,  
 11908 cx        the error indicator for the stream shall be set, *fgets()* shall return a null pointer, and shall set  
 11909        *errno* to indicate the error.

11910 **ERRORS**11911        Refer to *fgetc()*.11912 **EXAMPLES**11913        **Reading Input**

11914        The following example uses *fgets()* to read each line of input. {LINE\_MAX}, which defines the  
 11915        maximum size of the input line, is defined in the <limits.h> header.

```
11916 #include <stdio.h>
11917 ...
11918 char line[LINE_MAX];
11919 ...
11920 while (fgets(line, LINE_MAX, fp) != NULL) {
11921 ...
11922 }
11923 ...
```

11924 **APPLICATION USAGE**

11925        None.

11926 **RATIONALE**

11927        None.

11928 **FUTURE DIRECTIONS**

11929        None.

11930 **SEE ALSO**

11931 *fopen()*, *fread()*, *gets()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>

11932 **CHANGE HISTORY**

11933 First released in Issue 1. Derived from Issue 1 of the SVID.

11934 **Issue 4**

11935 In the DESCRIPTION, the text is changed to make it clear that the function reads bytes rather  
11936 than (possibly multi-byte) characters, and the list of functions that may cause the *st\_atime* field to  
11937 be updated is revised.

11938 **Issue 6**

11939 Extensions beyond the ISO C standard are now marked.

11940 The prototype for *fgets()* is changed for alignment with the ISO/IEC 9899:1999 standard.

## 11941 NAME

11942 fgetwc — get a wide-character code from a stream

## 11943 SYNOPSIS

11944 #include &lt;stdio.h&gt;

11945 #include &lt;wchar.h&gt;

11946 wint\_t fgetwc(FILE \*stream);

## 11947 DESCRIPTION

11948 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 11949 conflict between the requirements described here and the ISO C standard is unintentional. This  
 11950 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

11951 The *fgetwc()* function shall obtain the next character (if present) from the input stream pointed to  
 11952 by *stream*, convert that to the corresponding wide-character code, and advance the associated  
 11953 file position indicator for the stream (if defined).

11954 If an error occurs, the resulting value of the file position indicator for the stream is  
 11955 indeterminate.

11956 CX The *fgetwc()* function may mark the *st\_atime* field of the file associated with *stream* for update.  
 11957 The *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
 11958 *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns  
 11959 data not supplied by a prior call to *ungetc()* or *ungetwc()*.

## 11960 RETURN VALUE

11961 Upon successful completion, the *fgetwc()* function shall return the wide-character code of the  
 11962 character read from the input stream pointed to by *stream* converted to a type *wint\_t*. If the  
 11963 stream is at end-of-file, the end-of-file indicator for the stream shall be set and *fgetwc()* shall  
 11964 return WEOF. If a read error occurs, the error indicator for the stream shall be set, *fgetwc()* shall  
 11965 CX return WEOF, and shall set *errno* to indicate the error.

## 11966 ERRORS

11967 The *fgetwc()* function shall fail if data needs to be read and:

11968 CX [EAGAIN] The O\_NONBLOCK flag is set for the file descriptor underlying *stream* and the  
 11969 process would be delayed in the *fgetwc()* operation.

11970 CX [EBADF] The file descriptor underlying *stream* is not a valid file descriptor open for  
 11971 reading.

11972 CX [EINTR] The read operation was terminated due to the receipt of a signal, and no data  
 11973 was transferred.

11974 CX [EIO] A physical I/O error has occurred, or the process is in a background process  
 11975 group attempting to read from its controlling terminal, and either the process  
 11976 is ignoring or blocking the SIGTTIN signal or the process group is orphaned.  
 11977 This error may also be generated for implementation-defined reasons.

11978 CX [EOVERFLOW] The file is a regular file and an attempt was made to read at or beyond the  
 11979 offset maximum associated with the corresponding stream.

11980 The *fgetwc()* function may fail if:

11981 CX [ENOMEM] Insufficient storage space is available.

11982 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the  
 11983 capabilities of the device.

- 11984 CX [EILSEQ] The data obtained from the input stream does not form a valid character.
- 11985 **EXAMPLES**
- 11986 None.
- 11987 **APPLICATION USAGE**
- 11988 The *ferror()* or *feof()* functions must be used to distinguish between an error condition and an  
11989 end-of-file condition.
- 11990 **RATIONALE**
- 11991 None.
- 11992 **FUTURE DIRECTIONS**
- 11993 None.
- 11994 **SEE ALSO**
- 11995 *feof()*, *ferror()*, *fopen()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdio.h**>,  
11996 <**wchar.h**>
- 11997 **CHANGE HISTORY**
- 11998 First released in Issue 4. Derived from the MSE working draft.
- 11999 **Issue 4, Version 2**
- 12000 In the ERRORS section, the description of [EIO] is updated to include the case where a physical  
12001 I/O error occurs.
- 12002 **Issue 5**
- 12003 The Optional Header (OH) marking is removed from <**stdio.h**>.
- 12004 Large File Summit extensions are added.
- 12005 **Issue 6**
- 12006 Extensions beyond the ISO C standard are now marked.
- 12007 The following new requirements on POSIX implementations derive from alignment with the  
12008 Single UNIX Specification:
- 12009
  - The [EIO] and [Eoverflow] mandatory error conditions are added.
- 12010
  - The [ENOMEM], [ENXIO], and [EILSEQ] optional error conditions are added.

## 12011 NAME

12012 fgetws — get a wide-character string from a stream

## 12013 SYNOPSIS

12014 #include &lt;stdio.h&gt;

12015 #include &lt;wchar.h&gt;

12016 wchar\_t \*fgetws(wchar\_t \*restrict ws, int n,

12017 FILE \*restrict stream);

## 12018 DESCRIPTION

12019 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
12020 conflict between the requirements described here and the ISO C standard is unintentional. This  
12021 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

12022 The *fgetws()* function shall read characters from the *stream*, convert these to the corresponding  
12023 wide-character codes, place them in the **wchar\_t** array pointed to by *ws*, until *n*–1 characters are  
12024 read, or a <newline> character is read, converted, and transferred to *ws*, or an end-of-file  
12025 condition is encountered. The wide-character string, *ws*, is then terminated with a null wide-  
12026 character code.

12027 If an error occurs, the resulting value of the file position indicator for the stream is  
12028 indeterminate.

12029 CX The *fgetws()* function may mark the *st\_atime* field of the file associated with *stream* for update.  
12030 The *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
12031 *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns  
12032 data not supplied by a prior call to *ungetc()* or *ungetwc()*.

## 12033 RETURN VALUE

12034 Upon successful completion, *fgetws()* shall return *ws*. If the stream is at end-of-file, the end-of-  
12035 file indicator for the stream shall be set and *fgetws()* shall return a null pointer. If a read error  
12036 CX occurs, the error indicator for the stream shall be set, *fgetws()* shall return a null pointer, and  
12037 shall set *errno* to indicate the error.

## 12038 ERRORS

12039 Refer to *fgetwc()*.

## 12040 EXAMPLES

12041 None.

## 12042 APPLICATION USAGE

12043 None.

## 12044 RATIONALE

12045 None.

## 12046 FUTURE DIRECTIONS

12047 None.

## 12048 SEE ALSO

12049 *fopen()*, *fread()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>, <wchar.h>

## 12050 CHANGE HISTORY

12051 First released in Issue 4. Derived from the MSE working draft.

12052 **Issue 5**

12053 The Optional Header (OH) marking is removed from `<stdio.h>`.

12054 **Issue 6**

12055 Extensions beyond the ISO C standard are now marked.

12056 The prototype for `fgetws()` is changed for alignment with the ISO/IEC 9899:1999 standard.

12057 **NAME**

12058 fileno — map a stream pointer to a file descriptor

12059 **SYNOPSIS**

12060 #include <stdio.h>

12061 int fileno(FILE \*stream);

12062 **DESCRIPTION**

12063 The *fileno()* function shall return the integer file descriptor associated with the stream pointed to  
12064 by *stream*.

12065 **RETURN VALUE**

12066 Upon successful completion, *fileno()* shall return the integer value of the file descriptor  
12067 associated with *stream*. Otherwise, the value  $-1$  shall be returned and *errno* set to indicate the  
12068 error.

12069 **ERRORS**

12070 The *fileno()* function may fail if:

12071 [EBADF] The *stream* argument is not a valid stream.

12072 **EXAMPLES**

12073 None.

12074 **APPLICATION USAGE**

12075 None.

12076 **RATIONALE**

12077 Without some specification of which file descriptors are associated with these streams, it is  
12078 impossible for an application to set up the streams for another application it starts with *fork()*  
12079 and *exec*. In particular, it would not be possible to write a portable version of the *sh* command  
12080 interpreter (although there may be other constraints that would prevent that portability).

12081 **FUTURE DIRECTIONS**

12082 None.

12083 **SEE ALSO**

12084 *fdopen()*, *fopen()*, *stdin*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdio.h**>, Section  
12085 2.5.1 (on page 535)

12086 **CHANGE HISTORY**

12087 First released in Issue 1. Derived from Issue 1 of the SVID.

12088 **Issue 4**

12089 The [EBADF] error is marked as an extension.

12090 **Issue 6**

12091 The following new requirements on POSIX implementations derive from alignment with the  
12092 Single UNIX Specification:

- 12093 • The [EBADF] optional error condition is added.

12094 **NAME**

12095 flockfile, ftrylockfile, funlockfile — stdio locking functions

12096 **SYNOPSIS**

```
12097 TSF #include <stdio.h>
12098 void flockfile(FILE *file);
12099 int ftrylockfile(FILE *file);
12100 void funlockfile(FILE *file);
12101
```

12102 **DESCRIPTION**

12103 The *flockfile()*, *ftrylockfile()*, and *funlockfile()* functions provide for explicit application-level  
 12104 locking of stdio (**FILE\***) objects. These functions can be used by a thread to delineate a sequence  
 12105 of I/O statements that are executed as a unit.

12106 The *flockfile()* function is used by a thread to acquire ownership of a (**FILE\***) object.

12107 The *ftrylockfile()* function is used by a thread to acquire ownership of a (**FILE\***) object if the  
 12108 object is available; *ftrylockfile()* is a non-blocking version of *flockfile()*.

12109 The *funlockfile()* function is used to relinquish the ownership granted to the thread. The  
 12110 behavior is undefined if a thread other than the current owner calls the *funlockfile()* function.

12111 Logically, there is a lock count associated with each (**FILE\***) object. This count is implicitly  
 12112 initialized to zero when the (**FILE\***) object is created. The (**FILE\***) object is unlocked when the  
 12113 count is zero. When the count is positive, a single thread owns the (**FILE\***) object. When the  
 12114 *flockfile()* function is called, if the count is zero or if the count is positive and the caller owns the  
 12115 (**FILE\***) object, the count is incremented. Otherwise, the calling thread is suspended, waiting for  
 12116 the count to return to zero. Each call to *funlockfile()* decrements the count. This allows matching  
 12117 calls to *flockfile()* (or successful calls to *ftrylockfile()*) and *funlockfile()* to be nested.

12118 All functions that reference (**FILE\***) objects shall behave as if they use *flockfile()* and *funlockfile()*  
 12119 internally to obtain ownership of these (**FILE\***) objects.

12120 **RETURN VALUE**

12121 None for *flockfile()* and *funlockfile()*. The *ftrylockfile()* function shall return zero for success and  
 12122 non-zero to indicate that the lock cannot be acquired.

12123 **ERRORS**

12124 No errors are defined.

12125 **EXAMPLES**

12126 None.

12127 **APPLICATION USAGE**

12128 Applications using these functions may be subject to priority inversion, as discussed in the Base  
 12129 Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

12130 **RATIONALE**

12131 The *flockfile()* and *funlockfile()* functions provide an orthogonal mutual exclusion lock for each  
 12132 **FILE**. The *ftrylockfile()* function provides a non-blocking attempt to acquire a file lock,  
 12133 analogous to *pthread\_mutex\_trylock()*.

12134 These locks behave as if they are the same as those used internally by *stdio* for thread-safety.  
 12135 This both provides thread-safety of these functions without requiring a second level of internal  
 12136 locking and allows functions in *stdio* to be implemented in terms of other *stdio* functions.

12137 Application writers and implementors should be aware that there are potential deadlock  
 12138 problems on **FILE** objects. For example, the line-buffered flushing semantics of *stdio* (requested

12139 via {\_IOLBF}) require that certain input operations sometimes cause the buffered contents of  
12140 implementation-defined line-buffered output streams to be flushed. If two threads each hold the  
12141 lock on the other's **FILE**, deadlock ensues. This type of deadlock can be avoided by acquiring  
12142 **FILE** locks in a consistent order. In particular, the line-buffered output stream deadlock can  
12143 typically be avoided by acquiring locks on input streams before locks on output streams if a  
12144 thread would be acquiring both.

12145 In summary, threads sharing *stdio* streams with other threads can use *flockfile()* and *funlockfile()*  
12146 to cause sequences of I/O performed by a single thread to be kept bundled. The only case where  
12147 the use of *flockfile()* and *funlockfile()* is required is to provide a scope protecting uses of the  
12148 \*\_unlocked() functions/macros. This moves the cost/performance tradeoff to the optimal point.

#### 12149 **FUTURE DIRECTIONS**

12150 None.

#### 12151 **SEE ALSO**

12152 *getc\_unlocked()*, *putc\_unlocked()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdio.h**>

#### 12153 **CHANGE HISTORY**

12154 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

#### 12155 **Issue 6**

12156 These functions are marked as part of the Thread-Safe Functions option.

12157 **NAME**

12158 floor, floorf, floorl — floor function

12159 **SYNOPSIS**

12160 #include &lt;math.h&gt;

12161 double floor(double x);

12162 float floorf(float x);

12163 long double floorl(long double x);

12164 **DESCRIPTION**

12165 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 12166 conflict between the requirements described here and the ISO C standard is unintentional. This  
 12167 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

12168 These functions shall compute the largest integral value not greater than *x*.

12169 An application wishing to check for error situations should set *errno* to 0 before calling *floor()*. If  
 12170 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

12171 **RETURN VALUE**

12172 Upon successful completion, these functions shall return the largest integral value not greater  
 12173 than *x*, expressed as a **double**.

12174 **XSI** If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

12175 If the correct value would cause overflow, -HUGE\_VAL shall be returned and *errno* set to  
 12176 [ERANGE].

12177 **XSI** If *x* is ±Inf or ±0, the value of *x* shall be returned.12178 **ERRORS**

12179 These functions shall fail if:

12180 [ERANGE] The result would cause an overflow.

12181 These functions may fail if:

12182 **XSI** [EDOM] The value of *x* is NaN.12183 **XSI** No other errors shall occur.12184 **EXAMPLES**

12185 None.

12186 **APPLICATION USAGE**

12187 The integral value returned by *floor()* as a **double** might not be expressible as an **int** or **long**. The  
 12188 return value should be tested before assigning it to an integer type to avoid the undefined results  
 12189 of an integer overflow.

12190 The *floor()* function can only overflow when the floating point representation has  
 12191 DBL\_MANT\_DIG > DBL\_MAX\_EXP.

12192 **RATIONALE**

12193 None.

12194 **FUTURE DIRECTIONS**

12195 None.

12196 **SEE ALSO**

12197 *ceil()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

12198 **CHANGE HISTORY**

12199 First released in Issue 1. Derived from Issue 1 of the SVID.

12200 **Issue 4**

12201 References to *matherr()* are removed.

12202 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
12203 ISO C standard and to rationalize handling in the mathematics functions.

12204 The word **long** has been replaced with the words **long** in the APPLICATION USAGE section.

12205 The return value specified for [EDOM] is marked as an extension.

12206 **Issue 5**

12207 The DESCRIPTION is updated to indicate how an application should check for an error. This  
12208 text was previously published in the APPLICATION USAGE section.

12209 **Issue 6**

12210 The *floorf()* and *floorl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

12211 **NAME**

12212 fma, fmaf, fmal — floating-point multiply-add

12213 **SYNOPSIS**

12214 #include <math.h>

12215 double fma(double x, double y, double z);

12216 float fmaf(float x, float y, float z);

12217 long double fmal(long double x, long double y, long double z);

12218 **DESCRIPTION**

12219 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
12220 conflict between the requirements described here and the ISO C standard is unintentional. This  
12221 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

12222 These functions shall compute  $(x * y) + z$ , rounded as one ternary operation: they shall compute  
12223 the value (as if) to infinite precision and round once to the result format, according to the  
12224 rounding mode characterized by the value of FLT\_ROUNDS.

12225 An application wishing to check for error situations should set *errno* to 0 before calling these  
12226 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

12227 **RETURN VALUE**

12228 Upon successful completion, these functions shall return  $(x * y) + z$ , rounded as one ternary  
12229 operation.

12230 If *x*, *y*, or *z* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

12231 **ERRORS**

12232 These functions may fail if:

12233 [EDOM] The value of *x*, *y*, or *z* is NaN.

12234 **EXAMPLES**

12235 None.

12236 **APPLICATION USAGE**

12237 None.

12238 **RATIONALE**

12239 In many cases, clever use of floating (*fused*) multiply-add leads to much improved code; but its  
12240 unexpected use by the compiler can undermine carefully written code. The FP\_CONTRACT  
12241 macro can be used to disallow use of floating multiply-add; and the *fma()* function guarantees  
12242 its use where desired. Many current machines provide hardware floating multiply-add  
12243 instructions; software implementation can be used for others.

12244 **FUTURE DIRECTIONS**

12245 None.

12246 **SEE ALSO**

12247 The Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

12248 **CHANGE HISTORY**

12249 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

12250 **NAME**

12251 fmax, fmaxf, fmaxl — determine maximum numeric value of two floating-point numbers

12252 **SYNOPSIS**

12253 #include <math.h>

12254 double fmax(double x, double y);

12255 float fmaxf(float x, float y);

12256 long double fmaxl(long double x, long double y);

12257 **DESCRIPTION**

12258 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
12259 conflict between the requirements described here and the ISO C standard is unintentional. This  
12260 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

12261 These functions shall determine the maximum numeric value of their arguments. NaN  
12262 arguments shall be treated as missing data: if one argument is a NaN and the other numeric,  
12263 then the *fmax()*, *fmaxf()*, and *fmaxl()* functions shall choose the numeric value.

12264 An application wishing to check for error situations should set *errno* to 0 before calling these  
12265 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

12266 **RETURN VALUE**

12267 Upon successful completion, these functions shall return the maximum numeric value of their  
12268 arguments.

12269 If *x* and *y* are NaN, NaN shall be returned and *errno* may be set to [EDOM].

12270 **ERRORS**

12271 These functions may fail if:

12272 [EDOM] The value of *x* and *y* is NaN.

12273 **EXAMPLES**

12274 None.

12275 **APPLICATION USAGE**

12276 None.

12277 **RATIONALE**

12278 None.

12279 **FUTURE DIRECTIONS**

12280 None.

12281 **SEE ALSO**

12282 *fdim()*, *fmin()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

12283 **CHANGE HISTORY**

12284 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

12285 **NAME**

12286 `fmin`, `fminf`, `fminl` — determine minimum numeric value of two floating-point numbers

12287 **SYNOPSIS**

12288 `#include <math.h>`

12289 `double fmin(double x, double y);`

12290 `float fminf(float x, float y);`

12291 `long double fminl(long double x, long double y);`

12292 **DESCRIPTION**

12293 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
12294 conflict between the requirements described here and the ISO C standard is unintentional. This  
12295 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

12296 These functions shall determine the minimum numeric value of their arguments. NaN  
12297 arguments shall be treated as missing data: if one argument is a NaN and the other numeric,  
12298 then these functions shall choose the numeric value.

12299 An application wishing to check for error situations should set *errno* to 0 before calling these  
12300 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

12301 **RETURN VALUE**

12302 Upon successful completion, these functions shall return the minimum numeric value of their  
12303 arguments.

12304 If *x* and *y* are NaN, NaN shall be returned and *errno* may be set to [EDOM].

12305 **ERRORS**

12306 These functions may fail if:

12307 [EDOM] The value of *x* and *y* is NaN.

12308 **EXAMPLES**

12309 None.

12310 **APPLICATION USAGE**

12311 None.

12312 **RATIONALE**

12313 None.

12314 **FUTURE DIRECTIONS**

12315 None.

12316 **SEE ALSO**

12317 *fdim()*, *fmax()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<math.h>`

12318 **CHANGE HISTORY**

12319 First released in Issue 6. Derived from ISO/IEC 9899:1999 standard.

12320 **NAME**

12321 fmod, fmodf, fmodl — floating-point remainder value function

12322 **SYNOPSIS**

12323 #include &lt;math.h&gt;

12324 double fmod(double x, double y);

12325 float fmodf(float x, float y);

12326 long double fmodl(long double x, long double y);

12327 **DESCRIPTION**

12328 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 12329 conflict between the requirements described here and the ISO C standard is unintentional. This  
 12330 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

12331 These functions shall return the floating-point remainder of the division of  $x$  by  $y$ .

12332 An application wishing to check for error situations should set *errno* to 0 before calling *fmod()*. If  
 12333 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

12334 **RETURN VALUE**

12335 These functions shall return the value  $x-i*y$ , for some integer  $i$  such that, if  $y$  is non-zero, the  
 12336 result has the same sign as  $x$  and magnitude less than the magnitude of  $y$ .

12337 **XSI** If  $x$  or  $y$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].

12338 **XSI** If  $y$  is 0, NaN shall be returned and *errno* set to [EDOM], or 0 shall be returned and *errno* may be  
 12339 set to [EDOM].

12340 **XSI** If  $x$  is  $\pm\text{Inf}$ , either 0 shall be returned and *errno* set to [EDOM], or NaN shall be returned and *errno*  
 12341 may be set to [EDOM].

12342 If  $y$  is non-zero, *fmod*( $\pm 0, y$ ) shall return the value of  $x$ . If  $x$  is not  $\pm\text{Inf}$ , *fmod*( $x, \pm\text{Inf}$ ) shall return  
 12343 the value of  $x$ .

12344 If the result underflows, 0 shall be returned and *errno* may be set to [ERANGE].12345 **ERRORS**

12346 These functions may fail if:

12347 **XSI** [EDOM] One or both of the arguments is NaN, or  $y$  is 0, or  $x$  is  $\pm\text{Inf}$ .

12348 [ERANGE] The result underflows

12349 **XSI** No other errors shall occur.12350 **EXAMPLES**

12351 None.

12352 **APPLICATION USAGE**

12353 Portable applications should not call *fmod()* with  $y$  equal to 0, because the result is  
 12354 implementation-defined. The application should verify  $y$  is non-zero before calling *fmod()*.

12355 **RATIONALE**

12356 None.

12357 **FUTURE DIRECTIONS**

12358 None.

12359 **SEE ALSO**

12360 *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

12361 **CHANGE HISTORY**

12362 First released in Issue 1. Derived from Issue 1 of the SVID.

12363 **Issue 4**

12364 References to *matherr()* are removed.

12365 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
12366 ISO C standard and to rationalize error handling in the mathematics functions.

12367 The return value specified for [EDOM] is marked as an extension.

12368 **Issue 5**

12369 The DESCRIPTION is updated to indicate how an application should check for an error. This  
12370 text was previously published in the APPLICATION USAGE section.

12371 **Issue 6**

12372 The *fmodf()* and *fmodl()* functions are added for alignment with the ISO/IEC 9899:1999  
12373 standard.

12374 **NAME**

12375       fmtmsg — display a message in the specified format on standard error and/or a system console

12376 **SYNOPSIS**

```
12377 xSI #include <fmtmsg.h>
12378 int fmtmsg(long classification, const char *label, int severity,
12379 const char *text, const char *action, const char *tag);
12380
```

12381 **DESCRIPTION**12382       The *fmtmsg()* function can be used to display messages in a specified format instead of the  
12383 traditional *printf()* function.12384       Based on a message's classification component, *fmtmsg()* writes a formatted message either to  
12385 standard error, to the console, or to both.12386       A formatted message consists of up to five components as defined below. The component  
12387 *classification* is not part of a message displayed to the user, but defines the source of the message  
12388 and directs the display of the formatted message.

12389       *classification*       Contains identifiers from the following groups of major classifications and  
12390 subclassifications. Any one identifier from a subclass may be used in  
12391 combination with a single identifier from a different subclass. Two or more  
12392 identifiers from the same subclass should not be used together, with the  
12393 exception of identifiers from the display subclass. (Both display subclass  
12394 identifiers may be used so that messages can be displayed to both standard  
12395 error and the system console).

12396       **Major Classifications**12397           Identifies the source of the condition. Identifiers are: MM\_HARD  
12398 (hardware), MM\_SOFT (software), and MM\_FIRM (firmware).12399       **Message Source Subclassifications**12400           Identifies the type of software in which the problem is detected.  
12401 Identifiers are: MM\_APPL (application), MM\_UTIL (utility), and  
12402 MM\_OPSYS (operating system).12403       **Display Subclassifications**12404           Indicates where the message is to be displayed. Identifiers are:  
12405 MM\_PRINT to display the message on the standard error stream,  
12406 MM\_CONSOLE to display the message on the system console. One or  
12407 both identifiers may be used.12408       **Status Subclassifications**12409           Indicates whether the application can recover from the condition.  
12410 Identifiers are: MM\_RECOVER (recoverable) and MM\_NRECOV (non-  
12411 recoverable).12412       An additional identifier, MM\_NULLMC, indicates that no classification  
12413 component is supplied for the message.12414       *label*               Identifies the source of the message. The format is two fields separated by a  
12415 colon. The first field is up to 10 bytes, the second is up to 14 bytes.12416       *severity*           Indicates the seriousness of the condition. Identifiers for the levels of *severity*  
12417 are:

|       |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------|---------------------|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 12418 |                     | MM_HALT    | Indicates that the application has encountered a severe fault and is halting. Produces the string "HALT".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 12419 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12420 |                     | MM_ERROR   | Indicates that the application has detected a fault. Produces the string "ERROR".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 12421 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12422 |                     | MM_WARNING | Indicates a condition that is out of the ordinary, that might be a problem, and should be watched. Produces the string "WARNING".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 12423 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12424 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12425 |                     | MM_INFO    | Provides information about a condition that is not in error. Produces the string "INFO".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 12426 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12427 |                     | MM_NOSEV   | Indicates that no severity level is supplied for the message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 12428 | <i>text</i>         |            | Describes the error condition that produced the message. The character string is not limited to a specific size. If the character string is empty, then the text produced is unspecified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 12429 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12430 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12431 | <i>action</i>       |            | Describes the first step to be taken in the error-recovery process. The <i>fmtmsg()</i> function precedes the action string with the prefix: "TO FIX:". The <i>action</i> string is not limited to a specific size.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 12432 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12433 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12434 | <i>tag</i>          |            | An identifier that references on-line documentation for the message. Suggested usage is that <i>tag</i> includes the <i>label</i> and a unique identifying number. A sample <i>tag</i> is "XSI:cat:146".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 12435 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12436 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12437 |                     |            | The <i>MSGVERB</i> environment variable (for message verbosity) tells <i>fmtmsg()</i> which message components it is to select when writing messages to standard error. The value of <i>MSGVERB</i> is a colon-separated list of optional keywords. Valid <i>keywords</i> are: <i>label</i> , <i>severity</i> , <i>text</i> , <i>action</i> , and <i>tag</i> . If <i>MSGVERB</i> contains a keyword for a component and the component's value is not the component's null value, <i>fmtmsg()</i> includes that component in the message when writing the message to standard error. If <i>MSGVERB</i> does not include a keyword for a message component, that component is not included in the display of the message. The keywords may appear in any order. If <i>MSGVERB</i> is not defined, if its value is the null string, if its value is not of the correct format, or if it contains keywords other than the valid ones listed above, <i>fmtmsg()</i> selects all components. |
| 12438 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12439 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12440 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12441 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12442 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12443 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12444 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12445 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12446 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12447 |                     |            | <i>MSGVERB</i> affects only which components are selected for display to standard error. All message components are included in console messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 12448 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12449 | <b>RETURN VALUE</b> |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12450 |                     |            | The <i>fmtmsg()</i> function shall return one of the following values:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 12451 | MM_OK               |            | The function succeeded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| 12452 | MM_NOTOK            |            | The function failed completely.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12453 | MM_NOMSG            |            | The function was unable to generate a message on standard error, but otherwise succeeded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 12454 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12455 | MM_NOCON            |            | The function was unable to generate a console message, but otherwise succeeded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12456 |                     |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12457 | <b>ERRORS</b>       |            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 12458 |                     |            | None.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

12459 **EXAMPLES**

12460 1. The following example of *fmtmsg()*:

```
12461 fmtmsg(MM_PRINT, "XSI:cat", MM_ERROR, "illegal option",
12462 "refer to cat in user's reference manual", "XSI:cat:001")
```

12463 produces a complete message in the specified message format:

```
12464 XSI:cat: ERROR: illegal option
12465 TO FIX: refer to cat in user's reference manual XSI:cat:001
```

12466 2. When the environment variable *MSGVERB* is set as follows:

```
12467 MSGVERB=severity:text:action
```

12468 and Example 1 is used, *fmtmsg()* produces:

```
12469 ERROR: illegal option
12470 TO FIX: refer to cat in user's reference manual
```

12471 **APPLICATION USAGE**

12472 One or more message components may be systematically omitted from messages generated by  
12473 an application by using the null value of the argument for that component.

12474 **RATIONALE**

12475 None.

12476 **FUTURE DIRECTIONS**

12477 None.

12478 **SEE ALSO**

12479 *printf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <*fmtmsg.h*>

12480 **CHANGE HISTORY**

12481 First released in Issue 4, Version 2.

12482 **Issue 5**

12483 Moved from X/OPEN UNIX extension to BASE.

12484 **NAME**

12485       fnmatch — match a file name or a path name

12486 **SYNOPSIS**

12487       #include &lt;fnmatch.h&gt;

12488       int fnmatch(const char \**pattern*, const char \**string*, int *flags*);12489 **DESCRIPTION**

12490       The *fnmatch()* function shall match patterns as described in the Shell and Utilities volume of  
 12491       IEEE Std. 1003.1-200x, Section 2.14.1, Patterns Matching a Single Character, and Section 2.14.2,  
 12492       Patterns Matching Multiple Characters. It checks the string specified by the *string* argument to  
 12493       see if it matches the pattern specified by the *pattern* argument.

12494       The *flags* argument modifies the interpretation of *pattern* and *string*. It is the bitwise-inclusive OR  
 12495       of zero or more of the flags defined in <fnmatch.h>. If the FNM\_PATHNAME flag is set in *flags*,  
 12496       then a slash character ('/') in *string* shall be explicitly matched by a slash in *pattern*; it shall not  
 12497       be matched by either the asterisk or question-mark special characters, nor by a bracket  
 12498       expression. If the FNM\_PATHNAME flag is not set, the slash character is treated as an ordinary  
 12499       character.

12500       If FNM\_NOESCAPE is not set in *flags*, a backslash character ('\') in *pattern* followed by any  
 12501       other character shall match that second character in *string*. In particular, "\\\" shall match a  
 12502       backslash in *string*. If FNM\_NOESCAPE is set, a backslash character shall be treated as an  
 12503       ordinary character.

12504       If FNM\_PERIOD is set in *flags*, then a leading period ('.') in *string* shall match a period in  
 12505       *pattern*; as described by rule 2 in the Shell and Utilities volume of IEEE Std. 1003.1-200x, Section  
 12506       2.14.3, Patterns Used for File Name Expansion where the location of “leading” is indicated by  
 12507       the value of FNM\_PATHNAME:

- 12508       • If FNM\_PATHNAME is set, a period is “leading” if it is the first character in *string* or if it  
 12509       immediately follows a slash.
- 12510       • If FNM\_PATHNAME is not set, a period is “leading” only if it is the first character of *string*.

12511       If FNM\_PERIOD is not set, then no special restrictions are placed on matching a period.

12512 **RETURN VALUE**

12513       If *string* matches the pattern specified by *pattern*, then *fnmatch()* shall return 0. If there is no  
 12514       match, *fnmatch()* shall return FNM\_NOMATCH, which is defined in <fnmatch.h>. If an error  
 12515       occurs, *fnmatch()* shall return another non-zero value.

12516 **ERRORS**

12517       No errors are defined.

12518 **EXAMPLES**

12519       None.

12520 **APPLICATION USAGE**

12521       The *fnmatch()* function has two major uses. It could be used by an application or utility that  
 12522       needs to read a directory and apply a pattern against each entry. The *find* utility is an example of  
 12523       this. It can also be used by the *pax* utility to process its *pattern* operands, or by applications that  
 12524       need to match strings in a similar manner.

12525       The name *fnmatch()* is intended to imply *file name* match, rather than *path name* match. The  
 12526       default action of this function is to match file names, rather than path names, since it gives no  
 12527       special significance to the slash character. With the FNM\_PATHNAME flag, *fnmatch()* does  
 12528       match path names, but without tilde expansion, parameter expansion, or special treatment for a

12529 period at the beginning of a file name.

12530 **RATIONALE**

12531 This function replaced the REG\_FILENAME flag of *regcomp()* in early proposals of this volume  
12532 of IEEE Std. 1003.1-200x. It provides virtually the same functionality as the *regcomp()* and  
12533 *regexec()* functions using the REG\_FILENAME and REG\_FSLASH flags (the REG\_FSLASH flag  
12534 was proposed for *regcomp()*, and would have had the opposite effect from FNM\_PATHNAME),  
12535 but with a simpler function and less system overhead.

12536 **FUTURE DIRECTIONS**

12537 None.

12538 **SEE ALSO**

12539 *glob()*, *wordexp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <fnmatch.h>, the Shell  
12540 and Utilities volume of IEEE Std. 1003.1-200x

12541 **CHANGE HISTORY**

12542 First released in Issue 4. Derived from the ISO POSIX-2 standard.

12543 **Issue 5**

12544 Moved from POSIX2 C-language Binding to BASE.

## 12545 NAME

12546 fopen — open a stream

## 12547 SYNOPSIS

12548 #include &lt;stdio.h&gt;

12549 FILE \*fopen(const char \*restrict filename, const char \*restrict mode);

## 12550 DESCRIPTION

12551 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 12552 conflict between the requirements described here and the ISO C standard is unintentional. This  
 12553 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

12554 The *fopen()* function shall open the file whose path name is the string pointed to by *filename*, and  
 12555 associates a stream with it.

12556 The argument *mode* points to a string. If the string is one of the following, the file is open in the  
 12557 indicated mode. Otherwise, the behavior is undefined.

12558 *r* or *rb* Open file for reading.

12559 *w* or *wb* Truncate to zero length or create file for writing.

12560 *a* or *ab* Append; open or create file for writing at end-of-file.

12561 *r+* or *rb+* or *r+b* Open file for update (reading and writing).

12562 *w+* or *wb+* or *w+b* Truncate to zero length or create file for update.

12563 *a+* or *ab+* or *a+b* Append; open or create file for update, writing at end-of-file.

12564 CX The character 'b' has no effect, but is allowed for ISO C standard conformance. Opening a file  
 12565 with read mode (*r* as the first character in the *mode* argument) shall fail if the file does not exist or  
 12566 cannot be read.

12567 Opening a file with append mode (*a* as the first character in the *mode* argument) shall cause all  
 12568 subsequent writes to the file to be forced to the then current end-of-file, regardless of intervening  
 12569 calls to *fseek()*.

12570 When a file is opened with update mode ('+' as the second or third character in the *mode*  
 12571 argument), both input and output may be performed on the associated stream. However, the  
 12572 application shall ensure that output is not directly followed by input without an intervening call  
 12573 to *fflush()* or to a file positioning function (*fseek()*, *fsetpos()*, or *rewind()*), and input is not directly  
 12574 followed by output without an intervening call to a file positioning function, unless the input  
 12575 operation encounters end-of-file.

12576 When opened, a stream is fully buffered if and only if it can be determined not to refer to an  
 12577 interactive device. The error and end-of-file indicators for the stream shall be cleared.

12578 CX If *mode* is *w*, *a*, *w+*, or *a+*, and the file did not previously exist, upon successful completion,  
 12579 *fopen()* function shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime* fields of the file and  
 12580 the *st\_ctime* and *st\_mtime* fields of the parent directory.

12581 If *mode* is *w* or *w+* and the file did previously exist, upon successful completion, *fopen()* shall  
 12582 mark for update the *st\_ctime* and *st\_mtime* fields of the file. The *fopen()* function shall allocate a  
 12583 file descriptor as *open()* does.

12584 XSI After a successful call to the *fopen()* function, the orientation of the stream shall be cleared, the  
 12585 encoding rule shall be cleared, and the associated **mbstate\_t** object shall be set to describe an  
 12586 initial conversion state.

12587 The largest value that can be represented correctly in an object of type `off_t` shall be established  
 12588 as the offset maximum in the open file description.

#### 12589 RETURN VALUE

12590 Upon successful completion, `fopen()` shall return a pointer to the object controlling the stream.  
 12591 CX Otherwise, a null pointer shall be returned, and `errno` shall be set to indicate the error.

#### 12592 ERRORS

12593 The `fopen()` function shall fail if:

12594 CX [EACCES] Search permission is denied on a component of the path prefix, or the file  
 12595 exists and the permissions specified by `mode` are denied, or the file does not  
 12596 exist and write permission is denied for the parent directory of the file to be  
 12597 created.

12598 CX [EINTR] A signal was caught during `fopen()`.

12599 CX [EISDIR] The named file is a directory and `mode` requires write access.

12600 CX [ELOOP] A loop exists in symbolic links encountered during resolution of the `path`  
 12601 argument.

12602 CX [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

12603 CX [ENAMETOOLONG]

12604 The length of the `filename` argument exceeds {PATH\_MAX} or a path name  
 12605 component is longer than {NAME\_MAX}.

12606 CX [ENFILE] The maximum allowable number of files is currently open in the system.

12607 CX [ENOENT] A component of `filename` does not name an existing file or `filename` is an empty  
 12608 string.

12609 CX [ENOSPC] The directory or file system that would contain the new file cannot be  
 12610 expanded, the file does not exist, and it was to be created.

12611 CX [ENOTDIR] A component of the path prefix is not a directory.

12612 CX [ENXIO] The named file is a character special or block special file, and the device  
 12613 associated with this special file does not exist.

12614 CX [EOVERFLOW] The named file is a regular file and the size of the file cannot be represented  
 12615 correctly in an object of type `off_t`.

12616 CX [EROFS] The named file resides on a read-only file system and `mode` requires write  
 12617 access.

12618 The `fopen()` function may fail if:

12619 CX [EINVAL] The value of the `mode` argument is not valid.

12620 CX [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 12621 resolution of the `path` argument.

12622 CX [EMFILE] {FOPEN\_MAX} streams are currently open in the calling process.

12623 CX [EMFILE] {STREAM\_MAX} streams are currently open in the calling process.

12624 CX [ENAMETOOLONG]

12625 Path name resolution of a symbolic link produced an intermediate result  
 12626 whose length exceeds {PATH\_MAX}.

|       |    |           |                                                                                        |
|-------|----|-----------|----------------------------------------------------------------------------------------|
| 12627 | CX | [ENOMEM]  | Insufficient storage space is available.                                               |
| 12628 | CX | [ETXTBSY] | The file is a pure procedure (shared text) file that is being executed and <i>mode</i> |
| 12629 |    |           | requires write access.                                                                 |

12630 **EXAMPLES**12631 **Opening a File**

12632 The following example tries to open the file named **file** for reading. The *fopen()* function returns  
 12633 a file pointer that is used in subsequent *fgets()* and *fclose()* calls. If the program cannot open the  
 12634 file, it just ignores it.

```

12635 #include <stdio.h>
12636 ...
12637 FILE *fp;
12638 ...
12639 void rgrep(const char *file)
12640 {
12641 ...
12642 if ((fp = fopen(file, "r")) == NULL)
12643 return;
12644 ...
12645 }
```

12646 **APPLICATION USAGE**

12647 None.

12648 **RATIONALE**

12649 None.

12650 **FUTURE DIRECTIONS**

12651 None.

12652 **SEE ALSO**

12653 *fclose()*, *fdopen()*, *freopen()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<stdio.h>**

12654 **CHANGE HISTORY**

12655 First released in Issue 1. Derived from Issue 1 of the SVID.

12656 **Issue 4**

12657 In the DESCRIPTION, the descriptions of input and output operations on update streams are  
 12658 changed to be requirements on the application.

12659 The [EMFILE] error is added to the ERRORS section, and all the optional errors are marked as  
 12660 extensions.

12661 The following changes are incorporated for alignment with the ISO C standard:

- 12662 • The type of arguments *filename* and *mode* are changed from **char\*** to **const char\***.
- 12663 • In the DESCRIPTION, the use and settings of the *mode* argument are changed to support  
 12664 binary streams, and *setpos()* is added to the list of file positioning functions.

12665 The following change is incorporated for alignment with the FIPS requirements:

- 12666 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
 12667 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
 12668 an extension.

12669 **Issue 4, Version 2**

12670 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 12671 • It states that [ELOOP] is returned if too many symbolic links are encountered during path  
12672 name resolution.
- 12673 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an  
12674 intermediate result of path name resolution of a symbolic link.

12675 **Issue 5**

12676 Large File Summit extensions are added.

12677 **Issue 6**

12678 Extensions beyond the ISO C standard are now marked.

12679 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 12680 • The [ENAMETOOLONG] error is restored as an error dependent on \_POSIX\_NO\_TRUNC.  
12681 This is since behavior may vary from one file system to another.

12682 The following new requirements on POSIX implementations derive from alignment with the  
12683 Single UNIX Specification:

- 12684 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file  
12685 description. This change is to support large files.
- 12686 • In the ERRORS section, the [EOVERFLOW] condition is added. This change is to support  
12687 large files.
- 12688 • The [ELOOP] mandatory error condition is added.
- 12689 • The [EINVAL], [EMFILE], [ENAMETOOLONG], [ENOMEM], and [ETXTBSY] optional error  
12690 conditions are added.

12691 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

12692 The following changes are made for alignment with the ISO/IEC 9899: 1999 standard:

- 12693 • The prototype for *fopen()* is updated.
- 12694 • The DESCRIPTION is updated to note that if the argument *mode* points to a string other than  
12695 those listed, then the behavior is undefined.

12696 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
12697 [ELOOP] error condition is added.

## 12698 NAME

12699 fork — create a new process

## 12700 SYNOPSIS

12701 #include &lt;unistd.h&gt;

12702 pid\_t fork(void);

## 12703 DESCRIPTION

12704 The *fork()* function creates a new process. The new process (child process) shall be an exact copy  
 12705 of the calling process (parent process) except as detailed below:

- 12706 • The child process has a unique process ID.
- 12707 • The child process ID also does not match any active process group ID.
- 12708 • The child process has a different parent process ID (that is, the process ID of the parent  
 12709 process).
- 12710 • The child process has its own copy of the parent's file descriptors. Each of the child's file  
 12711 descriptors refers to the same open file description with the corresponding file descriptor of  
 12712 the parent.
- 12713 • The child process has its own copy of the parent's open directory streams. Each open  
 12714 directory stream in the child process may share directory stream positioning with the  
 12715 corresponding directory stream of the parent.
- 12716 XSI • The child process may have its own copy of the parent's message catalog descriptors.
- 12717 • The child process' values of *tms\_utime*, *tms\_stime*, *tms\_cutime*, and *tms\_cstime* are set to 0.
- 12718 • The time left until an alarm clock signal is reset to zero, and the alarm, if any, is canceled; see  
 12719 *alarm()*.
- 12720 XSI • All *semadj* values are cleared.
- 12721 • File locks set by the parent process are not inherited by the child process.
- 12722 • The set of signals pending for the child process is initialized to the empty set.
- 12723 XSI • Interval timers are reset in the child process.
- 12724 SEM • Any semaphores that are open in the parent process shall also be open in the child process.
- 12725 ML • The child process does not inherit any address space memory locks established by the parent  
 12726 process via calls to *mlockall()* or *mlock()*.
- 12727 MF|SHM • Memory mappings created in the parent are retained in the child process. MAP\_PRIVATE  
 12728 mappings inherited from the parent shall also be MAP\_PRIVATE mappings in the child, and  
 12729 any modifications to the data in these mappings made by the parent prior to calling *fork()*  
 12730 shall be visible to the child. Any modifications to the data in MAP\_PRIVATE mappings made  
 12731 by the parent after *fork()* returns shall be visible only to the parent. Modifications to the data  
 12732 in MAP\_PRIVATE mappings made by the child shall be visible only to the child.
- 12733 PS • For the SCHED\_FIFO and SCHED\_RR scheduling policies, the child process shall inherit the  
 12734 policy and priority settings of the parent process during a *fork()* function. For other  
 12735 scheduling policies, the policy and priority settings on *fork()* are implementation-defined.
- 12736 TMR • Per-process timers created by the parent are not inherited by the child process.
- 12737 MSG • The child process has its own copy of the message queue descriptors of the parent. Each of  
 12738 the message descriptors of the child refers to the same open message queue description as  
 12739 the corresponding message descriptor of the parent.

- 12740 AIO  
12741
- No asynchronous input or asynchronous output operations are inherited by the child process.
- 12742
- A process is created with a single thread. If a multi-threaded process calls *fork()*, the new process contains a replica of the calling thread and its entire address space, possibly including the states of mutexes and other resources. Consequently, to avoid errors, the child process may only execute async-signal-safe operations until such time as one of the *exec* functions is called. Fork handlers may be established by means of the *pthread\_atfork()* function in order to maintain application invariants across *fork()* calls.
- 12743  
12744  
12745  
12746 THR  
12747
- 12748 TRC TRI
- If the Trace option and the Trace Inherit option are both supported:
- 12749 If the calling process was being traced in a trace stream that had its inheritance policy set to  
12750 POSIX\_TRACE\_INHERITED, the child process shall be traced into that trace stream, and the  
12751 child process shall inherit the parent's mapping of trace event names to trace event type  
12752 identifiers. If the trace stream in which the calling process was being traced had its  
12753 inheritance policy set to POSIX\_TRACE\_CLOSE\_FOR\_CHILD, the child process shall not be  
12754 traced into that trace stream. The inheritance policy is set by a call to the  
12755 *posix\_trace\_attr\_setinherited()* function.
- 12756 TRC
- If the Trace option is supported, but the Trace Inherit option is not supported:
- 12757 The child process shall not be traced into any of the trace streams of its parent process.
- 12758
- If the Trace option is supported, the child process of a trace controller process shall not  
12759 control the trace streams controlled by its parent process.
- 12760 CPT The initial value of the CPU-time clock of the child process shall be set to zero.
- 12761 TCT The initial value of the CPU-time clock of the single thread of the child process shall be set to  
12762 zero.

### 12763 **Notes to Reviewers**

- 12764 *This section with side shading will not appear in the final copy. - Ed.*
- 12765 Check this text is correct after new addenda are rolled in. (Ref D1, XSH, ERN 126)
- 12766 For process characteristics not defined by this volume of IEEE Std. 1003.1-200x, their inheritance
- 12767 is not defined by this volume of IEEE Std. 1003.1-200x. Process characteristics defined by this
- 12768 volume of IEEE Std. 1003.1-200x have their inheritance explicitly defined.
- 12769 After *fork()*, both the parent and the child processes are capable of executing independently
- 12770 before either one terminates.

### 12771 **RETURN VALUE**

- 12772 Upon successful completion, *fork()* shall return 0 to the child process and shall return the
- 12773 process ID of the child process to the parent process. Both processes shall continue to execute
- 12774 from the *fork()* function. Otherwise, -1 shall be returned to the parent process, no child process
- 12775 shall be created, and *errno* shall be set to indicate the error.

### 12776 **ERRORS**

- 12777 The *fork()* function shall fail if:
- 12778 [EAGAIN] The system lacked the necessary resources to create another process, or the
- 12779 system-imposed limit on the total number of processes under execution
- 12780 system-wide or by a single user {CHILD\_MAX} would be exceeded.
- 12781 The *fork()* function may fail if:

12782 [ENOMEM] Insufficient storage space is available.

12783 **EXAMPLES**

12784 None.

12785 **APPLICATION USAGE**

12786 None.

12787 **RATIONALE**

12788 Many historical implementations have timing windows where a signal sent to a process group  
 12789 (for example, an interactive SIGINT) just prior to or during execution of *fork()* is delivered to the  
 12790 parent following the *fork()* but not to the child because the *fork()* code clears the child's set of  
 12791 pending signals. This volume of IEEE Std. 1003.1-200x does not require, or even permit, this  
 12792 behavior. However, it is pragmatic to expect that problems of this nature may continue to exist  
 12793 in implementations that appear to conform to this volume of IEEE Std. 1003.1-200x and pass  
 12794 available verification suites. This behavior is only a consequence of the implementation failing to  
 12795 make the interval between signal generation and delivery totally invisible. From the  
 12796 application's perspective, a *fork()* call should appear atomic. A signal that is generated prior to  
 12797 the *fork()* should be delivered prior to the *fork()*. A signal sent to the process group after the  
 12798 *fork()* should be delivered to both parent and child. The implementation may actually initialize  
 12799 internal data structures corresponding to the child's set of pending signals to include signals  
 12800 sent to the process group during the *fork()*. Since the *fork()* call can be considered as atomic  
 12801 from the application's perspective, the set would be initialized as empty and such signals would  
 12802 have arrived after the *fork()*; see also <signal.h>.

12803 One approach that has been suggested to address the problem of signal inheritance across *fork()*  
 12804 is to add an [EINTR] error, which would be returned when a signal is detected during the call.  
 12805 While this is preferable to losing signals, it was not considered an optimal solution. Although it  
 12806 is not recommended for this purpose, such an error would be an allowable extension for an  
 12807 implementation.

12808 The [ENOMEM] error value is reserved for those implementations that detect and distinguish  
 12809 such a condition. This condition occurs when an implementation detects that there is not enough  
 12810 memory to create the process. This is intended to be returned when [EAGAIN] is inappropriate  
 12811 because there can never be enough memory (either primary or secondary storage) to perform the  
 12812 operation. Because *fork()* duplicates an existing process, this must be a condition where there is  
 12813 sufficient memory for one such process, but not for two. Many historical implementations  
 12814 actually return [ENOMEM] due to temporary lack of memory, a case that is not generally  
 12815 distinct from [EAGAIN] from the perspective of a portable application.

12816 Part of the reason for including the optional error [ENOMEM] is because the SVID specifies it  
 12817 and it should be reserved for the error condition specified there. The condition is not applicable  
 12818 on many implementations.

12819 IEEE Std. 1003.1-1988 neglected to require concurrent execution of the parent and child of *fork()*.  
 12820 A system that single-threads processes was clearly not intended and is considered an  
 12821 unacceptable "toy implementation" of this volume of IEEE Std. 1003.1-200x. The only objection  
 12822 anticipated to the phrase "executing independently" is testability, but this assertion should be  
 12823 testable. Such tests require that both the parent and child can block on a detectable action of the  
 12824 other, such as a write to a pipe or a signal. An interactive exchange of such actions should be  
 12825 possible for the system to conform to the intent of this volume of IEEE Std. 1003.1-200x.

12826 The [EAGAIN] error exists to warn applications that such a condition might occur. Whether it  
 12827 occurs or not is not in any practical sense under the control of the application because the  
 12828 condition is usually a consequence of the user's use of the system, not of the application's code.  
 12829 Thus, no application can or should rely upon its occurrence under any circumstances, nor

12830 should the exact semantics of what concept of “user” is used be of concern to the application  
12831 writer. Validation writers should be cognizant of this limitation.

12832 There are two reasons why POSIX programmers call *fork()*. One reason is to create a new thread  
12833 of control within the same program (which was originally only possible in POSIX by creating a  
12834 new process); the other is to create a new process running a different program. In the latter case,  
12835 the call to *fork()* is soon followed by a call to one of the *exec* functions.

12836 The general problem with making *fork()* work in a multi-threaded world is what to do with all  
12837 of the threads. There are two alternatives. One is to copy all of the threads into the new process.  
12838 This causes the programmer or implementation to deal with threads that are suspended on  
12839 system calls or that might be about to execute system calls that should not be executed in the  
12840 new process. The other alternative is to copy only the thread that calls *fork()*. This creates the  
12841 difficulty that the state of process-local resources is usually held in process memory. If a thread  
12842 that is not calling *fork()* holds a resource, that resource is never released in the child process  
12843 because the thread whose job it is to release the resource does not exist in the child process.

12844 When a programmer is writing a multi-threaded program, the first described use of *fork()*,  
12845 creating new threads in the same program, is provided by the *pthread\_create()* function. The  
12846 *fork()* function is thus used only to run new programs, and the effects of calling functions that  
12847 require certain resources between the call to *fork()* and the call to an *exec* function are undefined.

12848 The addition of the *forkall()* function to the standard was considered and rejected. The *forkall()*  
12849 function lets all the threads in the parent be duplicated in the child. This essentially duplicates  
12850 the state of the parent in the child. This allows threads in the child to continue processing and  
12851 allows locks and the state to be preserved without explicit *pthread\_atfork()* code. The calling  
12852 process has to ensure that the threads processing state that is shared between the parent and  
12853 child (that is, file descriptors or MAP\_SHARED memory) behaves properly after *forkall()*. For  
12854 example, if a thread is reading a file descriptor in the parent when *forkall()* is called, then two  
12855 threads (one in the parent and one in the child) are reading the file descriptor after the *forkall()*.  
12856 If this is not desired behavior, the parent process has to synchronize with such threads before  
12857 calling *forkall()*.

12858 When *forkall()* is called, threads, other than the calling thread, that are in POSIX System  
12859 Interfaces functions that can return with an [EINTR] error may have those functions return  
12860 [EINTR] if the implementation cannot ensure that the function behaves correctly in the parent  
12861 and child. In particular, *pthread\_cond\_wait()* and *pthread\_cond\_timedwait()* need to return in order  
12862 to ensure that the condition has not changed. These functions can be awakened by a spurious  
12863 condition wakeup rather than returning [EINTR].

#### 12864 FUTURE DIRECTIONS

12865 None.

#### 12866 SEE ALSO

12867 *alarm()*, *exec*, *fcntl()*, *posix\_trace\_attr\_getinherited()*, *posix\_trace\_trid\_eventid\_open()*, *semop()*,  
12868 *signal()*, *times()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<sys/types.h>`,  
12869 `<unistd.h>`

#### 12870 CHANGE HISTORY

12871 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 12872 Issue 4

12873 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on  
12874 XSI-conformant systems.

12875 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 12876           • The argument list is explicitly defined as **void**.
- 12877           • Though functionally identical to Issue 3, the DESCRIPTION has been reorganized to improve  
12878           clarity and to align more closely with the ISO POSIX-1 standard.
- 12879           • The description of the [EAGAIN] error is updated to indicate that this error can also be  
12880           returned if a system lacks the resources to create another process.
- 12881 **Issue 4, Version 2**
- 12882           The DESCRIPTION is changed for X/OPEN UNIX conformance to identify that interval timers  
12883           are reset in the child process.
- 12884 **Issue 5**
- 12885           The DESCRIPTION is changed for alignment with the POSIX Realtime Extension and the POSIX  
12886           Threads Extension.
- 12887 **Issue 6**
- 12888           The following new requirements on POSIX implementations derive from alignment with the  
12889           Single UNIX Specification:
- 12890           • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
12891           required for conforming implementations of previous POSIX specifications, it was not  
12892           required for UNIX applications.
- 12893           The following changes were made to align with the IEEE P1003.1a draft standard:
- 12894           • The effect of `fork()` on a pending alarm call in the child process is clarified.
- 12895           The description of CPU-time clock semantics is added for alignment with  
12896           IEEE Std. 1003.1d-1999.
- 12897           The description of tracing semantics is added for alignment with IEEE Std. 1003.1q-2000.

12898 **NAME**

12899 `fpathconf`, `pathconf` — get configurable path name variables

12900 **SYNOPSIS**

12901 `#include <unistd.h>`

12902 `long fpathconf(int fildev, int name);`

12903 `long pathconf(const char *path, int name);`

12904 **DESCRIPTION**

12905 The `fpathconf()` and `pathconf()` functions provide a method for the application to determine the  
 12906 current value of a configurable limit or option (*variable*) that is associated with a file or directory.

12907 For `pathconf()`, the *path* argument points to the path name of a file or directory.

12908 For `fpathconf()`, the *fildev* argument is an open file descriptor.

12909 The *name* argument represents the variable to be queried relative to that file or directory.  
 12910 Implementations shall support all of the variables listed in the following table and may support  
 12911 others. The variables in the following table come from `<limits.h>` or `<unistd.h>` and the  
 12912 symbolic constants, defined in `<unistd.h>`, are the corresponding values used for *name*. Support  
 12913 for some path name configuration variables is dependent on implementation options (see  
 12914 shading and margin codes in the table below). Where an implementation option is not  
 12915 supported, the variable need not be supported.

12916

12917

|           | Variable                   | Value of <i>name</i>   | Notes |
|-----------|----------------------------|------------------------|-------|
| 12918     | {FILESIZEBITS}             | _PC_FILESIZEBITS       | 3, 4  |
| 12919     | {LINK_MAX}                 | _PC_LINK_MAX           | 1     |
| 12920     | {MAX_CANON}                | _PC_MAX_CANON          | 2     |
| 12921     | {MAX_INPUT}                | _PC_MAX_INPUT          | 2     |
| 12922     | {NAME_MAX}                 | _PC_NAME_MAX           | 3, 4  |
| 12923     | {PATH_MAX}                 | _PC_PATH_MAX           | 4, 5  |
| 12924     | {PIPE_BUF}                 | _PC_PIPE_BUF           | 6     |
| 12925 ADV | {POSIX_ALLOC_SIZE_MIN}     | _PC_ALLOC_SIZE_MIN     |       |
| 12926 ADV | {POSIX_REC_INCR_XFER_SIZE} | _PC_REC_INCR_XFER_SIZE |       |
| 12927 ADV | {POSIX_REC_MAX_XFER_SIZE}  | _PC_REC_MAX_XFER_SIZE  |       |
| 12928 ADV | {POSIX_REC_MIN_XFER_SIZE}  | _PC_REC_MIN_XFER_SIZE  |       |
| 12929 ADV | {POSIX_REC_XFER_ALIGN}     | _PC_REC_XFER_ALIGN     |       |
| 12930     | {SYMLINK_MAX}              | _PC_SYMLINK_MAX        | 4, 9  |
| 12931     | _POSIX_CHOWN_RESTRICTED    | _PC_CHOWN_RESTRICTED   | 7     |
| 12932     | _POSIX_NO_TRUNC            | _PC_NO_TRUNC           | 3, 4  |
| 12933     | _POSIX_VDISABLE            | _PC_VDISABLE           | 2     |
| 12934     | _POSIX_ASYNC_IO            | _PC_ASYNC_IO           | 8     |
| 12935     | _POSIX_PRIO_IO             | _PC_PRIO_IO            | 8     |
| 12936     | _POSIX_SYNC_IO             | _PC_SYNC_IO            | 8     |

12937 **Notes:**

- 12938 1. If *path* or *fildev* refers to a directory, the value returned applies to the directory  
 12939 itself.
- 12940 2. If *path* or *fildev* does not refer to a terminal file, it is unspecified whether an  
 12941 implementation supports an association of the variable name with the specified  
 12942 file.

- 12943 3. If *path* or *filde*s refers to a directory, the value returned applies to file names  
12944 within the directory.
- 12945 4. If *path* or *filde*s does not refer to a directory, it is unspecified whether an  
12946 implementation supports an association of the variable name with the specified  
12947 file.
- 12948 5. If *path* or *filde*s refers to a directory, the value returned is the maximum length  
12949 of a relative path name when the specified directory is the working directory.
- 12950 6. If *path* refers to a FIFO, or *filde*s refers to a pipe or FIFO, the value returned  
12951 applies to the referenced object. If *path* or *filde*s refers to a directory, the value  
12952 returned applies to any FIFO that exists or can be created within the directory.  
12953 If *path* or *filde*s refers to any other type of file, it is unspecified whether an  
12954 implementation supports an association of the variable name with the specified  
12955 file.
- 12956 7. If *path* or *filde*s refers to a directory, the value returned applies to any files, other  
12957 than directories, that exist or can be created within the directory.
- 12958 8. If *path* or *filde*s refers to a directory, it is unspecified whether an implementation  
12959 supports an association of the variable name with the specified file.
- 12960 9. If *path* or *filde*s refers to a directory, the value returned is the maximum length  
12961 of the string that a symbolic link in that directory can contain.

**12962 RETURN VALUE**

12963 If *name* is an invalid value, both *pathconf()* and *fpathconf()* shall return  $-1$  and set *errno* to  
12964 indicate the error.

12965 If the variable corresponding to *name* has no limit for the *path* or file descriptor, both *pathconf()*  
12966 and *fpathconf()* shall return  $-1$  without changing *errno*. If the implementation needs to use *path*  
12967 to determine the value of *name* and the implementation does not support the association of *name*  
12968 with the file specified by *path*, or if the process did not have appropriate privileges to query the  
12969 file specified by *path*, or *path* does not exist, *pathconf()* shall return  $-1$  and set *errno* to indicate the  
12970 error.

12971 If the implementation needs to use *filde*s to determine the value of *name* and the implementation  
12972 does not support the association of *name* with the file specified by *filde*s, or if *filde*s is an invalid  
12973 file descriptor, *fpathconf()* shall return  $-1$  and set *errno* to indicate the error.

12974 Otherwise, *pathconf()* or *fpathconf()* shall return the current variable value for the file or  
12975 directory without changing *errno*. The value returned shall not be more restrictive than the  
12976 corresponding value available to the application when it was compiled with the  
12977 implementation's `<limits.h>` or `<unistd.h>`.

**12978 ERRORS**

12979 The *pathconf()* function shall fail if:

- |       |          |                                                                                  |  |
|-------|----------|----------------------------------------------------------------------------------|--|
| 12980 | [EINVAL] | The value of <i>name</i> is not valid.                                           |  |
| 12981 | [ELOOP]  | A loop exists in symbolic links encountered during resolution of the <i>path</i> |  |
| 12982 |          | argument.                                                                        |  |

12983 The *pathconf()* function may fail if:

- |       |          |                                                                                     |  |
|-------|----------|-------------------------------------------------------------------------------------|--|
| 12984 | [EACCES] | Search permission is denied for a component of the path prefix.                     |  |
| 12985 | [EINVAL] | The implementation does not support an association of the variable <i>name</i> with |  |
| 12986 |          | the specified file.                                                                 |  |

|       |                          |                                                                                                                                                                                                                                                                                                                                                                    |
|-------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 12987 | [ELOOP]                  | More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.                                                                                                                                                                                                                                                             |
| 12988 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 12989 | [ENAMETOOLONG]           |                                                                                                                                                                                                                                                                                                                                                                    |
| 12990 |                          | The length of the <i>path</i> argument exceeds {PATH_MAX} or a path name component is longer than {NAME_MAX}.                                                                                                                                                                                                                                                      |
| 12991 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 12992 | [ENAMETOOLONG]           |                                                                                                                                                                                                                                                                                                                                                                    |
| 12993 |                          | As a result of encountering a symbolic link in resolution of the <i>path</i> argument, the length of the substituted path name string exceeded {PATH_MAX}.                                                                                                                                                                                                         |
| 12994 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 12995 | [ENOENT]                 | A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.                                                                                                                                                                                                                                                                       |
| 12996 | [ENOTDIR]                | A component of the path prefix is not a directory.                                                                                                                                                                                                                                                                                                                 |
| 12997 |                          | The <i>fpathconf()</i> function shall fail if:                                                                                                                                                                                                                                                                                                                     |
| 12998 | [EINVAL]                 | The value of <i>name</i> is not valid.                                                                                                                                                                                                                                                                                                                             |
| 12999 |                          | The <i>fpathconf()</i> function may fail if:                                                                                                                                                                                                                                                                                                                       |
| 13000 | [EBADF]                  | The <i>fdes</i> argument is not a valid file descriptor.                                                                                                                                                                                                                                                                                                           |
| 13001 | [EINVAL]                 | The implementation does not support an association of the variable <i>name</i> with the specified file.                                                                                                                                                                                                                                                            |
| 13002 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13003 | <b>EXAMPLES</b>          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13004 |                          | None.                                                                                                                                                                                                                                                                                                                                                              |
| 13005 | <b>APPLICATION USAGE</b> |                                                                                                                                                                                                                                                                                                                                                                    |
| 13006 |                          | None.                                                                                                                                                                                                                                                                                                                                                              |
| 13007 | <b>RATIONALE</b>         |                                                                                                                                                                                                                                                                                                                                                                    |
| 13008 |                          | The <i>pathconf()</i> function was proposed immediately after the <i>sysconf()</i> function when it was realized that some configurable values may differ across file system, directory, or device boundaries.                                                                                                                                                     |
| 13009 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13010 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13011 |                          | For example, {NAME_MAX} frequently changes between System V and BSD-based file systems; System V uses a maximum of 14, BSD 255. On an implementation that provides both types of file systems, an application would be forced to limit all path name components to 14 bytes, as this would be the value specified in <limits.h> on such a system.                  |
| 13012 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13013 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13014 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13015 |                          | Therefore, various useful values can be queried on any path name or file descriptor, assuming that the appropriate permissions are in place.                                                                                                                                                                                                                       |
| 13016 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13017 |                          | The value returned for the variable {PATH_MAX} indicates the longest relative path name that could be given if the specified directory is the process' current working directory. A process may not always be able to generate a name that long and use it if a subdirectory in the path name crosses into a more restrictive file system.                         |
| 13018 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13019 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13020 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13021 |                          | The value returned for the variable _POSIX_CHOWN_RESTRICTED also applies to directories that do not have file systems mounted on them. The value may change when crossing a mount point, so applications that need to know should check for each directory. (An even easier check is to try the <i>chown()</i> function and look for an error in case it happens.) |
| 13022 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13023 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13024 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13025 |                          | Unlike the values returned by <i>sysconf()</i> , the path name-oriented variables are potentially more volatile and are not guaranteed to remain constant throughout the process' lifetime. For example, in between two calls to <i>pathconf()</i> , the file system in question may have been unmounted and remounted with different characteristics.             |
| 13026 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13027 |                          |                                                                                                                                                                                                                                                                                                                                                                    |
| 13028 |                          |                                                                                                                                                                                                                                                                                                                                                                    |

13029 Also note that most of the errors are optional. If one of the variables always has the same value  
 13030 on an implementation, the implementation need not look at *path* or *filde*s to return that value and  
 13031 is, therefore, not required to detect any of the errors except the meaning of [EINVAL] that  
 13032 indicates that the value of *name* is not valid for that variable.

13033 If the value of any of the limits are indeterminate (logically infinite), they are defined in  
 13034 `<limits.h>` and the *pathconf()* and *fpathconf()* functions return `-1` without changing *errno*. This  
 13035 can be distinguished from the case of giving an unrecognized *name* argument because *errno* is set  
 13036 to [EINVAL] in this case.

13037 Since `-1` is a valid return value for the *pathconf()* and *fpathconf()* functions, applications should  
 13038 set *errno* to zero before calling them and check *errno* only if the return value is `-1`.

13039 For the case of {SYMLINK\_MAX}, since both *pathconf()* and *open()* follow symbolic links, there  
 13040 is no way that *path* or *filde*s could refer to a symbolic link.

#### 13041 FUTURE DIRECTIONS

13042 None.

#### 13043 SEE ALSO

13044 *confstr()*, *sysconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<limits.h>`, `<unistd.h>`,  
 13045 the Shell and Utilities volume of IEEE Std. 1003.1-200x

#### 13046 CHANGE HISTORY

13047 First released in Issue 3.

13048 Entry included for alignment with the POSIX.1-1988 standard.

#### 13049 Issue 4

13050 The *fpathconf()* function now has the full **long** return type in the SYNOPSIS section.

13051 The following changes have been made for alignment with the ISO POSIX-1 standard:

- 13052 • The type of argument *path* is changed from **char\*** to **const char\***. Also, the return value of  
 13053 both functions is changed from **long** to **long**.
- 13054 • In the DESCRIPTION, the words “The behavior is undefined if” have been replaced by “it is  
 13055 unspecified whether an implementation supports an association of the variable name with  
 13056 the specified file” in notes 2, 4, and 6.
- 13057 • In the RETURN VALUE section, errors associated with the use of *path* and *filde*s, when an  
 13058 implementation does not support the requested association, are now specified separately.
- 13059 • The requirement that *errno* be set to indicate the error is added.

13060 The following change is incorporated for alignment with the FIPS requirements:

- 13061 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
 13062 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
 13063 an extension.

#### 13064 Issue 4, Version 2

13065 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 13066 • It states that [ELOOP] is returned if too many symbolic links are encountered during path  
 13067 name resolution.
- 13068 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an  
 13069 intermediate result of path name resolution of a symbolic link.

13070 **Issue 5**

13071 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.

13072 Large File Summit extensions are added.

13073 **Issue 6**

13074 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 13075 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.
- 13076 This is since behavior may vary from one file system to another.

13077 The following new requirements on POSIX implementations derive from alignment with the  
13078 Single UNIX Specification:

- 13079 • The DESCRIPTION is updated to include {FILESIZEBITS}.
- 13080 • The [ELOOP] mandatory error condition is added.
- 13081 • A second [ENAMETOOLONG] is added as an optional error condition.

13082 The following changes were made to align with the IEEE P1003.1a draft standard:

- 13083 • The `_PC_SYMLINK_MAX` entry is added to the table in the DESCRIPTION.

13084 The `pathconf()` variables {`POSIX_ALLOC_SIZE_MIN`}, {`POSIX_REC_INCR_XFER_SIZE`},  
13085 {`POSIX_REC_MAX_XFER_SIZE`}, {`POSIX_REC_MIN_XFER_SIZE`},  
13086 {`POSIX_REC_XFER_ALIGN`} and their associated names are added for alignment with  
13087 IEEE Std. 1003.1d-1999.

13088 **NAME**

13089 fpclassify — classify real floating type

13090 **SYNOPSIS**

13091 #include <math.h>

13092 int fpclassify(real-floating x);

13093 **DESCRIPTION**

13094 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
13095 conflict between the requirements described here and the ISO C standard is unintentional. This  
13096 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

13097 The *fpclassify()* macro shall classify its argument value as NaN, infinite, normal, subnormal,  
13098 zero, or into another implementation-defined category. First, an argument represented in a  
13099 format wider than its semantic type is converted to its semantic type. Then classification is based  
13100 on the type of the argument.

13101 **RETURN VALUE**

13102 The *fpclassify()* macro shall return the value of the number classification macro appropriate to  
13103 the value of its argument.

13104 **ERRORS**

13105 No errors are defined.

13106 **EXAMPLES**

13107 None.

13108 **APPLICATION USAGE**

13109 None.

13110 **RATIONALE**

13111 None.

13112 **FUTURE DIRECTIONS**

13113 None.

13114 **SEE ALSO**

13115 *isfinite()*, *isinf()*, *isnan()*, *isnormal()*, *signbit()*, the Base Definitions volume of  
13116 IEEE Std. 1003.1-200x, <math.h>

13117 **CHANGE HISTORY**

13118 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

## 13119 NAME

13120 fprintf, printf, snprintf, sprintf — print formatted output

## 13121 SYNOPSIS

13122 #include &lt;stdio.h&gt;

13123 int fprintf(FILE \*restrict *stream*, const char \*restrict *format*, ...);13124 int printf(const char \*restrict *format*, ...);13125 int snprintf(char \*restrict *s*, size\_t *n*, const char \*restrict *format*, ...);13126 int sprintf(char \*restrict *s*, const char \*restrict *format*, ...);

## 13127 DESCRIPTION

13128 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 13129 conflict between the requirements described here and the ISO C standard is unintentional. This  
 13130 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

13131 The *fprintf()* function places output on the named output *stream*. The *printf()* function places  
 13132 output on the standard output stream *stdout*. The *sprintf()* function places output followed by  
 13133 the null byte, '\0', in consecutive bytes starting at *s*; it is the user's responsibility to ensure that  
 13134 enough space is available.

13135 XSI The *snprintf()* function is identical to *sprintf()* with the addition of the *n* argument, which states  
 13136 the size of the buffer referred to by *s*. If *n* is greater than zero, but not large enough to hold all  
 13137 output bytes specified by the format, output bytes beyond the *n*-1st are discarded rather than  
 13138 being written into the array.

13139 If copying takes place between objects that overlap as a result of a call to *sprintf()* or *snprintf()*,  
 13140 the results are undefined.

13141 Each of these functions converts, formats, and prints its arguments under control of the *format*.  
 13142 The *format* is a character string, beginning and ending in its initial shift state, if any. The *format* is  
 13143 composed of zero or more directives: *ordinary characters*, which are simply copied to the output  
 13144 stream, and *conversion specifications*, each of which results in the fetching of zero or more  
 13145 arguments. The results are undefined if there are insufficient arguments for the *format*. If the  
 13146 *format* is exhausted while arguments remain, the excess arguments are evaluated but are  
 13147 otherwise ignored.

13148 XSI Conversions can be applied to the *n*th argument after the *format* in the argument list, rather than  
 13149 to the next unused argument. In this case, the conversion character '%' (see below) is replaced  
 13150 by the sequence "%n\$", where *n* is a decimal integer in the range [1, {NL\_ARGMAX}], giving the  
 13151 position of the argument in the argument list. This feature provides for the definition of format  
 13152 strings that select arguments in an order appropriate to specific languages (see the EXAMPLES  
 13153 section).

13154 In format strings containing the "%n\$" form of conversion specifications, numbered arguments  
 13155 in the argument list can be referenced from the format string as many times as required.

13156 In format strings containing the '%' form of conversion specifications, each argument in the  
 13157 argument list is used exactly once.

13158 All forms of the *fprintf()* functions allow for the insertion of a language-dependent radix  
 13159 character in the output string. The radix character is defined in the program's locale (category  
 13160 LC\_NUMERIC). In the POSIX locale, or in a locale where the radix character is not defined, the  
 13161 radix character defaults to a period ('.').

13162 XSI Each conversion specification is introduced by the '%' character or by the character sequence  
 13163 "%n\$", after which the following appear in sequence:

- 13164           • Zero or more *flags* (in any order), which modify the meaning of the conversion specification.
  - 13165           • An optional minimum *field width*. If the converted value has fewer bytes than the field
  - 13166           width, it shall be padded with spaces by default on the left; it shall be padded on the right, if
  - 13167           the left-adjustment flag ('-'), described below, is given to the field width. The field width
  - 13168           takes the form of an asterisk ('\*'), described below, or a decimal integer.
  - 13169           • An optional *precision* that gives the minimum number of digits to appear for the *d*, *i*, *o*, *u*, *x*,
  - 13170           and *X* conversions; the number of digits to appear after the radix character for the *e*, *E*, and *f*
  - 13171           conversions; the maximum number of significant digits for the *g* and *G* conversions; or the
  - 13172 XSI       maximum number of bytes to be printed from a string in *s* and *S* conversions. The precision
  - 13173           takes the form of a period ('.') followed either by an asterisk ('\*'), described below, or an
  - 13174           optional decimal digit string, where a null digit string is treated as 0. If a precision appears
  - 13175           with any other conversion character, the behavior is undefined.
  - 13176           • An optional length modifier that specifies the size of the argument.
  - 13177           • A *conversion character* that indicates the type of conversion to be applied.
- 13178       A field width, or precision, or both, may be indicated by an asterisk ('\*'). In this case an
- 13179       argument of type **int** supplies the field width or precision. Applications shall ensure that
- 13180       arguments specifying field width, or precision, or both appear in that order before the argument,
- 13181       if any, to be converted. A negative field width is taken as a '-' flag followed by a positive field
- 13182 XSI       width. A negative precision is taken as if the precision were omitted. In format strings
- 13183       containing the "%n\$" form of a conversion specification, a field width or precision may be
- 13184       indicated by the sequence "\*m\$", where *m* is a decimal integer in the range [1,{NL\_ARGMAX}]
- 13185       giving the position in the argument list (after the format argument) of an integer argument
- 13186       containing the field width or precision, for example:
- ```
13187       printf("%1$d:%2$.*3$d:%4$.*3$d\n", hour, min, precision, sec);
```
- 13188 The *format* can contain either numbered argument specifications (that is, "%n\$" and "*m\$"), or
- 13189 unnumbered argument specifications (that is, '%' and '*'), but normally not both. The only
- 13190 exception to this is that "%%" can be mixed with the "%n\$" form. The results of mixing
- 13191 numbered and unnumbered argument specifications in a *format* string are undefined. When
- 13192 numbered argument specifications are used, specifying the *N*th argument requires that all the
- 13193 leading arguments, from the first to the (*N*-1)th, are specified in the format string.
- 13194 The flag characters and their meanings are:
- 13195 XSI ' The integer portion of the result of a decimal conversion (%i, %d, %u, %f, %g, or %G)
 - 13196 shall be formatted with thousands' grouping characters. For other conversions the
 - 13197 behavior is undefined. The non-monetary grouping character is used.
 - 13198 - The result of the conversion shall be left-justified within the field. The conversion is
 - 13199 right-justified if this flag is not specified.
 - 13200 + The result of a signed conversion shall always begin with a sign ('+' or '-'). The
 - 13201 conversion shall begin with a sign only when a negative value is converted if this flag is
 - 13202 not specified.
 - 13203 <space> If the first character of a signed conversion is not a sign or if a signed conversion results
 - 13204 in no characters, a space shall be prefixed to the result. This means that if the space and
 - 13205 '+' flags both appear, the space flag shall be ignored.
 - 13206 # This flag specifies that the value is to be converted to an alternative form. For *o*
 - 13207 conversion, it increases the precision (if necessary) to force the first digit of the result to
 - 13208 be 0. For *x* or *X* conversions, a non-zero result shall have 0x (or 0X) prefixed to it. For *e*,
 - 13209 *E*, *f*, *g*, or *G* conversions, the result shall always contain a radix character, even if no

13210 digits follow the radix character. Without this flag, a radix character appears in the
 13211 result of these conversions only if a digit follows it. For *g* and *G* conversions, trailing
 13212 zeros shall *not* be removed from the result as they normally are. For other conversions,
 13213 the behavior is undefined.

13214 **0** For *d*, *i*, *o*, *u*, *x*, *X*, *e*, *E*, *f*, *g*, and *G* conversions, leading zeros (following any indication of
 13215 sign or base) are used to pad to the field width; no space padding is performed. If the
 13216 '0' and '-' flags both appear, the '0' flag is ignored. For *d*, *i*, *o*, *u*, *x*, and *X*
 13217 XSJ conversions, if a precision is specified, the '0' flag is ignored. If the '0' and '\\'
 13218 flags both appear, the grouping characters are inserted before zero padding. For other
 13219 conversions, the behavior is undefined.

13220 The length modifiers and their meanings are:

13221 **hh** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **signed char**
 13222 or **unsigned char** argument (the argument will have been promoted according to the
 13223 integer promotions, but its value shall be converted to **signed char** or **unsigned char**
 13224 before printing); or that a following *n* conversion specifier applies to a pointer to a
 13225 **signed char** argument.

13226 **h** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **short** or
 13227 **unsigned short** argument (the argument will have been promoted according to the
 13228 integer promotions, but its value shall be converted to **short** or **unsigned short** before
 13229 printing); or that a following *n* conversion specifier applies to a pointer to a **short**
 13230 argument.

13231 **l** (ell) Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **long** or
 13232 **unsigned long** argument; that a following *n* conversion specifier applies to a pointer to
 13233 a **long** argument; that a following *c* conversion specifier applies to a **wint_t** argument;
 13234 that a following *s* conversion specifier applies to a pointer to a **wchar_t** argument; or
 13235 has no effect on a following *a*, *A*, *e*, *E*, *f*, *F*, *g*, or *G* conversion specifier.

13236 **ll** (ell-ell) Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **long long** or
 13237 **unsigned long long** argument; or that a following *n* conversion specifier applies to a
 13238 pointer to a **long long** argument.

13239 **j** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to an **intmax_t** or
 13240 **uintmax_t** argument; or that a following *n* conversion specifier applies to a pointer to
 13241 an **intmax_t** argument.

13242 **z** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **size_t** or the
 13243 corresponding signed integer type argument; or that a following *n* conversion specifier
 13244 applies to a pointer to a signed integer type corresponding to **size_t** argument.

13245 **t** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **ptrdiff_t** or
 13246 the corresponding **unsigned** type argument; or that a following *n* conversion specifier
 13247 applies to a pointer to a **ptrdiff_t** argument.

13248 **L** Specifies that a following *a*, *A*, *e*, *E*, *f*, *F*, *g*, or *G* conversion specifier applies to a **long**
 13249 **double** argument.

13250 If a length modifier appears with any conversion specifier other than as specified above, the
 13251 behavior is undefined.

13252 The conversion characters and their meanings are:

13253 **d, i** The **int** argument is converted to a signed decimal in the style `[-]dddd`. The precision
 13254 specifies the minimum number of digits to appear; if the value being converted can be
 13255 represented in fewer digits, it shall be expanded with leading zeros. The default

13256		precision is 1. The result of converting 0 with an explicit precision of 0 is no characters.
13257	<i>o</i>	The unsigned argument is converted to unsigned octal format in the style <i>ddd</i> . The
13258		precision specifies the minimum number of digits to appear; if the value being
13259		converted can be represented in fewer digits, it shall be expanded with leading zeros.
13260		The default precision is 1. The result of converting 0 with an explicit precision of 0 is no
13261		characters.
13262	<i>u</i>	The unsigned argument is converted to unsigned decimal format in the style <i>ddd</i> . The
13263		precision specifies the minimum number of digits to appear; if the value being
13264		converted can be represented in fewer digits, it shall be expanded with leading zeros.
13265		The default precision is 1. The result of converting 0 with an explicit precision of 0 is no
13266		characters.
13267	<i>x</i>	The unsigned argument is converted to unsigned hexadecimal format in the style <i>ddd</i> ;
13268		the letters "abcdef" are used. The precision specifies the minimum number of digits
13269		to appear; if the value being converted can be represented in fewer digits, it shall be
13270		expanded with leading zeros. The default precision is 1. The result of converting 0 with
13271		an explicit precision of 0 is no characters.
13272	<i>X</i>	Behaves the same as the <i>x</i> conversion character except that letters "ABCDEF" are used
13273		instead of "abcdef".
13274	<i>f, F</i>	The double argument is converted to decimal notation in the style <i>[-]ddd.ddd</i> , where
13275		the number of digits after the radix character is equal to the precision specification. If
13276		the precision is missing, it is taken as 6; if the precision is explicitly 0 and no '#' flag is
13277		present, no radix character appears. If a radix character appears, at least one digit
13278		appears before it. The value is rounded to the appropriate number of digits.
13279		A double argument representing an infinity is converted in one of the styles <i>[-]inf</i> or
13280		<i>[-]infinity</i> ; which style is implementation-defined. A double argument representing a
13281		NaN is converted in one of the styles <i>[-]nan</i> or <i>[-]nan(n-char-sequence)</i> ; which style, and
13282		the meaning of any <i>n-char-sequence</i> , is implementation-defined. The <i>F</i> conversion
13283		specifier produces INF, INFINITY, or NAN instead of inf, infinity, or nan, respectively.
13284	<i>e, E</i>	The double argument is converted in the style <i>[-]d.ddde±dd</i> , where there is one digit
13285		before the radix character (which is non-zero if the argument is non-zero) and the
13286		number of digits after it is equal to the precision; if the precision is missing, it is taken
13287		as 6; if the precision is 0 and no '#' flag is present, no radix character appears. The
13288		value is rounded to the appropriate number of digits. The <i>E</i> conversion character will
13289		produce a number with <i>E</i> instead of <i>e</i> introducing the exponent. The exponent always
13290		contains at least two digits. If the value is 0, the exponent is 0.
13291		A double argument representing an infinity or NaN is converted in the style of an <i>f</i> or <i>F</i>
13292		conversion specifier.
13293	<i>g, G</i>	The double argument is converted in the style <i>f</i> or <i>e</i> (or in the style <i>E</i> in the case of a <i>G</i>
13294		conversion character), with the precision specifying the number of significant digits. If
13295		an explicit precision is 0, it is taken as 1. The style used depends on the value
13296		converted; style <i>e</i> (or <i>E</i>) shall be used only if the exponent resulting from such a
13297		conversion is less than -4 or greater than or equal to the precision. Trailing zeros are
13298		removed from the fractional portion of the result; a radix character appears only if it is
13299		followed by a digit.
13300		A double argument representing an infinity or NaN is converted in the style of an <i>f</i> or <i>F</i>
13301		conversion specifier.

13302	a, A	A double argument representing a floating-point number is converted in the style <code>[-]0xh.hhhh pld</code> , where there is one hexadecimal digit (which is non-zero if the argument is a normalized floating-point number and is otherwise unspecified) before the decimal-point character 235) and the number of hexadecimal digits after it is equal to the precision; if the precision is missing and <code>FLT_RADIX</code> is a power of 2, then the precision is sufficient for an exact representation of the value; if the precision is missing and <code>FLT_RADIX</code> is not a power of 2, then the precision is sufficient to distinguish 236) values of type double , except that trailing zeros may be omitted; if the precision is zero and the '#' flag is not specified, no decimal-point character appears. The letters "abcdef" are used for a conversion and the letters "ABCDEF" for A conversion. The A conversion specifier produces a number with 'X' and 'P' instead of 'x' and 'p'. The exponent always contains at least one digit, and only as many more digits as necessary to represent the decimal exponent of 2. If the value is zero, the exponent is zero.
13303		
13304		
13305		
13306		
13307		
13308		
13309		
13310		
13311		
13312		A double argument representing an infinity or NaN is converted in the style of an <i>f</i> or <i>F</i> conversion specifier.
13317		
13318	c	The int argument is converted to an unsigned char , and the resulting byte is written. If an <i>l</i> (ell) qualifier is present, the wint_t argument is converted as if by an <i>ls</i> conversion specification with no precision and an argument that points to a two-element array of type wchar_t , the first element of which contains the wint_t argument to the <i>ls</i> conversion specification and the second element contains a null wide character.
13319		
13320		
13321		
13322		
13323		
13324		
13325		
13326		
13327		
13328	s	The argument shall be a pointer to an array of char . Bytes from the array are written up to (but not including) any terminating null byte. If the precision is specified, no more than that many bytes shall be written. If the precision is not specified or is greater than the size of the array, the application shall ensure that the array contains a null byte. If an <i>l</i> (ell) qualifier is present, the argument shall be a pointer to an array of type wchar_t . Wide characters from the array are converted to characters (each as if by a call to the <i>wcrtomb()</i> function, with the conversion state described by an mbstate_t object initialized to zero before the first wide character is converted) up to and including a terminating null wide character. The resulting characters shall be written up to (but not including) the terminating null character (byte). If no precision is specified, the application shall ensure that the array contains a null wide character. If a precision is specified, no more than that many characters (bytes) shall be written (including shift sequences, if any), and the array shall contain a null wide character if, to equal the character sequence length given by the precision, the function would need to access a wide character one past the end of the array. In no case is a partial character written.
13329		
13330		
13331		
13332		
13333		
13334		
13335		
13336		
13337		
13338	p	The argument shall be a pointer to void . The value of the pointer is converted to a sequence of printable characters, in an implementation-defined manner.
13341		
13342	n	The argument shall be a pointer to an integer into which is written the number of bytes written to the output so far by this call to one of the <i>fprintf()</i> functions. No argument is converted.
13343		
13344		
13345	XSI	C Same as <i>lc</i> .
13346	XSI	S Same as <i>ls</i> .
13347	%	Print a '%'; no argument is converted. The entire conversion specification shall be "%%".
13348		

13349 If a conversion specification does not match one of the above forms, the behavior is undefined. If
 13350 any argument is not the correct type for the corresponding conversion specification, the
 13351 behavior is undefined.

13352 In no case does a nonexistent or small field width cause truncation of a field; if the result of a
 13353 conversion is wider than the field width, the field is simply expanded to contain the conversion
 13354 result. Characters generated by *fprintf()* and *printf()* are printed as if *fputc()* had been called.

13355 For *a* and *A* conversions, if `FLT_RADIX` is a power of 2, the value is correctly rounded to a
 13356 hexadecimal floating number with the given precision.

13357 If `FLT_RADIX` is not a power of 2, the result should be one of the two adjacent numbers in
 13358 hexadecimal floating style with the given precision, with the extra stipulation that the error
 13359 should have a correct sign for the current rounding direction.

13360 For *e*, *E*, *f*, *F*, *g*, and *G* conversions, if the number of significant decimal digits is at most
 13361 `DECIMAL_DIG`, then the result should be correctly rounded. If the number of significant
 13362 decimal digits is more than `DECIMAL_DIG` but the source value is exactly representable with
 13363 `DECIMAL_DIG` digits, then the result should be an exact representation with trailing zeros.
 13364 Otherwise, the source value is bounded by two adjacent decimal strings "*L* < *U*", both having
 13365 `DECIMAL_DIG` significant digits; the value of the resultant decimal string "*D*" should satisfy "*L*
 13366 <= *D* <= *U*", with the extra stipulation that the error should have a correct sign for the current
 13367 rounding direction.

13368 CX The *st_ctime* and *st_mtime* fields of the file shall be marked for update between the call to a
 13369 successful execution of *fprintf()* or *printf()* and the next successful completion of a call to *fflush()*
 13370 or *fclose()* on the same stream or a call to *exit()* or *abort()*.

13371 RETURN VALUE

13372 Upon successful completion, the *fprintf()* and *printf()* functions shall return the number of bytes
 13373 transmitted.

13374 Upon successful completion, the *sprintf()* function shall return the number of bytes written to *s*,
 13375 excluding the terminating null byte.

13376 Upon successful completion, the *snprintf()* function shall return the number of bytes that would
 13377 be written to *s* had *n* been sufficiently large excluding the terminating null byte.

13378 If an output error was encountered, these functions shall return a negative value.

13379 If the value of *n* is zero on a call to *snprintf()*, nothing shall be written, the number of bytes that
 13380 would have been written had *n* been sufficiently large excluding the terminating null shall be
 13381 returned, and *s* may be a null pointer.

13382 ERRORS

13383 For the conditions under which *fprintf()* and *printf()* fail and may fail, refer to *fputc()* or
 13384 *fputwc()*.

13385 In addition, all forms of *fprintf()* may fail if:

13386 XSI [EILSEQ] A wide-character code that does not correspond to a valid character has been
 13387 detected.

13388 XSI [EINVAL] There are insufficient arguments.

13389 The *printf()* and *fprintf()* functions may fail if:

13390 XSI [ENOMEM] Insufficient storage space is available.

13391 The *snprintf()* function shall fail if:

13392 XSI [Eoverflow] The value of *n* is greater than {INT_MAX} or the number of bytes needed to
 13393 hold the output excluding the terminating null is greater than {INT_MAX}.

13394 EXAMPLES

13395 **Printing Language-Independent Date and Time**

13396 The following statement can be used to print date and time using a language-independent
 13397 format:

```
13398 printf(format, weekday, month, day, hour, min);
```

13399 For American usage, *format* could be a pointer to the following string:

```
13400 "%s, %s %d, %d:%.2d\n"
```

13401 This example would produce the following message:

```
13402 Sunday, July 3, 10:02
```

13403 For German usage, *format* could be a pointer to the following string:

```
13404 "%1$s, %3$d. %2$s, %4$d:%5$.2d\n"
```

13405 This definition of *format* would produce the following message:

```
13406 Sonntag, 3. Juli, 10:02
```

13407 **Printing File Information**

13408 The following example prints information about the type, permissions, and number of links of a
 13409 specific file in a directory.

13410 The first two calls to *printf()* use data decoded from a previous *stat()* call. The user-defined
 13411 *strperm()* function shall return a string similar to the one at the beginning of the output for the
 13412 following command:

```
13413 ls -l
```

13414 The next call to *printf()* outputs the owner's name if it is found using *getpwuid()*; the *getpwuid()*
 13415 function shall return a **passwd** structure from which the name of the user is extracted. If the user
 13416 name is not found, the program instead prints out the numeric value of the user ID.

13417 The next call prints out the group name if it is found using *getgrgid()*; *getgrgid()* is very similar to
 13418 *getpwuid()* except that it shall return group information based on the group number. Once
 13419 again, if the group is not found, the program prints the numeric value of the group for the entry.

13420 The final call to *printf()* prints the size of the file.

```
13421 #include <stdio.h>
```

```
13422 #include <sys/types.h>
```

```
13423 #include <pwd.h>
```

```
13424 #include <grp.h>
```

```
13425 char *strperm (mode_t);
```

```
13426 ...
```

```
13427 struct stat statbuf;
```

```
13428 struct passwd *pwd;
```

```
13429 struct group *grp;
```

```
13430 ...
```

```
13431 printf("%10.10s", strperm (statbuf.st_mode));
```

```

13432     printf("%4d", statbuf.st_nlink);
13433     if ((pwd = getpwuid(statbuf.st_uid)) != NULL)
13434         printf(" %-8.8s", pwd->pw_name);
13435     else
13436         printf(" %-8ld", (long) statbuf.st_uid);
13437     if ((grp = getgrgid(statbuf.st_gid)) != NULL)
13438         printf(" %-8.8s", grp->gr_name);
13439     else
13440         printf(" %-8ld", (long) statbuf.st_gid);
13441     printf("%9jd", (intmax_t) statbuf.st_size);
13442     ...

```

13443 **Printing a Localized Date String**

13444 The following example gets a localized date string. The *nl_langinfo()* function shall return the
 13445 localized date string, which specifies the order and layout of the date. The *strftime()* function
 13446 takes this information and, using the **tm** structure for values, places the date and time
 13447 information into *datestring*. The *printf()* function then outputs *datestring* and the name of the
 13448 entry.

```

13449     #include <stdio.h>
13450     #include <time.h>
13451     #include <langinfo.h>
13452     ...
13453     struct dirent *dp;
13454     struct tm *tm;
13455     char datestring[256];
13456     ...
13457     strftime(datestring, sizeof(datestring), nl_langinfo (D_T_FMT), tm);
13458     printf(" %s %s\n", datestring, dp->d_name);
13459     ...

```

13460 **Printing Error Information**

13461 The following example uses *fprintf()* to write error information to standard error.

13462 In the first group of calls, the program tries to open the password lock file named **LOCKFILE**. If
 13463 the file already exists, this is an error, as indicated by the **O_EXCL** flag on the *open()* function. If
 13464 the call fails, the program assumes that someone else is updating the password file, and the
 13465 program exits.

13466 The next group of calls saves a new password file as the current password file by creating a link
 13467 between **LOCKFILE** and the new password file **PASSWDFILE**.

```

13468     #include <sys/types.h>
13469     #include <sys/stat.h>
13470     #include <fcntl.h>
13471     #include <stdio.h>
13472     #include <stdlib.h>
13473     #include <unistd.h>
13474     #include <string.h>
13475     #include <errno.h>

```

```

13476     #define LOCKFILE "/etc/ptmp"
13477     #define PASSWDFILE "/etc/passwd"
13478     ...
13479     int pfd;
13480     ...
13481     if ((pfd = open(LOCKFILE, O_WRONLY | O_CREAT | O_EXCL,
13482                  S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH)) == -1)
13483     {
13484         fprintf(stderr, "Cannot open /etc/ptmp. Try again later.\n");
13485         exit(1);
13486     }
13487     ...
13488     if (link(LOCKFILE, PASSWDFILE) == -1) {
13489         fprintf(stderr, "Link error: %s\n", strerror(errno));
13490         exit(1);
13491     }
13492     ...

```

13493 **Printing Usage Information**

13494 The following example checks to make sure the program has the necessary arguments, and uses
 13495 *fprintf()* to print usage information if the expected number of arguments is not present.

```

13496     #include <stdio.h>
13497     #include <stdlib.h>
13498     ...
13499     char *Options = "hdbt1";
13500     ...
13501     if (argc < 2) {
13502         fprintf(stderr, "Usage: %s -%s <file\n", argv[0], Options); exit(1);
13503     }
13504     ...

```

13505 **Formatting a Decimal String**

13506 The following example prints a key and data pair on *stdout*. Note use of the '*' (asterisk) in the
 13507 format string; this ensures the correct number of decimal places for the element based on the
 13508 number of elements requested.

```

13509     #include <stdio.h>
13510     ...
13511     long i;
13512     char *keyst;
13513     int elementlen, len;
13514     ...
13515     while (len < elementlen) {
13516         ...
13517         printf("%s Element%0*ld\n", keyst, elementlen, i);
13518         ...
13519     }

```

13520 **Creating a File Name**

13521 The following example creates a file name using information from a previous *getpwnam()*
 13522 function that returned the HOME directory of the user.

```
13523 #include <stdio.h>
13524 #include <sys/types.h>
13525 #include <unistd.h>
13526 ...
13527 char filename[PATH_MAX+1];
13528 struct passwd *pw;
13529 ...
13530 sprintf(filename, "%s/%d.out", pw->pw_dir, getpid());
13531 ...
```

13532 **Reporting an Event**

13533 The following example loops until an event has timed out. The *pause()* function waits forever
 13534 unless it receives a signal. The *fprintf()* statement should never occur due to the possible return
 13535 values of *pause()*.

```
13536 #include <stdio.h>
13537 #include <unistd.h>
13538 #include <string.h>
13539 #include <errno.h>
13540 ...
13541 while (!event_complete) {
13542     ...
13543     if (pause() != -1 || errno != EINTR)
13544         fprintf(stderr, "pause: unknown error: %s\n", strerror(errno));
13545 }
13546 ...
```

13547 **Printing Monetary Information**

13548 The following example uses *strfmon()* to convert a number and store it as a formatted monetary
 13549 string named *convbuf*. If the first number is printed, the program prints the format and the
 13550 description; otherwise, it just prints the number.

```
13551 #include <monetary.h>
13552 #include <stdio.h>
13553 ...
13554 struct tblfmt {
13555     char *format;
13556     char *description;
13557 };
13558 struct tblfmt table[] = {
13559     { "%n", "default formatting" },
13560     { "%11n", "right align within an 11 character field" },
13561     { "%#5n", "aligned columns for values up to 99,999" },
13562     { "%=*#5n", "specify a fill character" },
13563     { "%=0#5n", "fill characters do not use grouping" },
13564     { "%^#5n", "disable the grouping separator" },
13565     { "%^#5.0n", "round off to whole units" },
```

```

13566         { "%^#5.4n", "increase the precision" },
13567         { "%(#5n", "use an alternative pos/neg style" },
13568         { "%!(#5n", "disable the currency symbol" },
13569     };
13570     ...
13571     float input[3];
13572     int i, j;
13573     char convbuf[100];
13574     ...
13575     strfmon(convbuf, sizeof(convbuf), table[i].format, input[j]);
13576
13577     if (j == 0) {
13578         printf("%s%s%s\n", table[i].format,
13579             convbuf, table[i].description);
13580     }
13581     else {
13582         printf("%s\n", convbuf);
13583     }
13584     ...

```

13584 APPLICATION USAGE

13585 If the application calling *fprintf()* has any objects of type **wint_t** or **wchar_t**, it must also include
 13586 the **<wchar.h>** header to have these objects defined.

13587 RATIONALE

13588 None.

13589 FUTURE DIRECTIONS

13590 None.

13591 SEE ALSO

13592 *fputc()*, *fscanf()*, *setlocale()*, *wcrtomb()*, the Base Definitions volume of IEEE Std. 1003.1-200x,
 13593 **<stdio.h>**, **<wchar.h>**, the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

13594 CHANGE HISTORY

13595 First released in Issue 1. Derived from Issue 1 of the SVID.

13596 Issue 4

13597 In the DESCRIPTION, references to *langinfo* data are marked as extensions. The reference to
 13598 *langinfo* data is removed from the description of the radix character.

13599 The ' ' (single-quote) flag is added to the list of flag characters and marked as an extension.
 13600 This flag directs that numeric conversion is formatted with the decimal grouping character.

13601 The detailed description of these functions is provided here instead of under *printf()*.

13602 The information in the APPLICATION USAGE section is moved to the DESCRIPTION. A new
 13603 APPLICATION USAGE section is added.

13604 The [EILSEQ] error is added to the ERRORS section and all errors are marked as extensions.

13605 The following changes are incorporated for alignment with the ISO C standard:

- 13606 • The type of the *format* arguments is changed from **char*** to **const char***.
- 13607 • The DESCRIPTION is reworded or presented differently in a number of places for alignment
 13608 with the ISO C standard, and also for clarity. There are no functional changes, except as
 13609 noted elsewhere in this CHANGE HISTORY section.

- 13610 The following changes are incorporated for alignment with the MSE working draft:
- 13611 • The *C* and *S* conversion characters are added, indicating respectively a wide character of type
 - 13612 `wchar_t` and pointer to a wide-character string of type `wchar_t*` in the argument list.
- 13613 **Issue 4, Version 2**
- 13614 The [ENOMEM] error is added to the ERRORS section as an optional error.
- 13615 **Issue 5**
- 13616 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the *l* (ell) qualifier can
- 13617 now be used with *c* and *s* conversion characters.
- 13618 The `snprintf()` function is new in Issue 5.
- 13619 **Issue 6**
- 13620 Extensions beyond the ISO C standard are now marked.
- 13621 The description of `snprintf()` has been aligned with The Open Group Base Resolution bwg98-006.
- 13622 This aligns `snprintf()` with historic BSD behavior.
- 13623 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 13624 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:
- 13625 • The prototypes for `fprintf()`, `printf()`, `snprintf()`, and `sprintf()` are updated, and the XSI
 - 13626 shading is removed from `snprintf()`.
 - 13627 • The DESCRIPTION is updated.

13628 **NAME**

13629 fputc — put a byte on a stream

13630 **SYNOPSIS**

13631 #include <stdio.h>

13632 int fputc(int *c*, FILE **stream*);13633 **DESCRIPTION**

13634 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 13635 conflict between the requirements described here and the ISO C standard is unintentional. This
 13636 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

13637 The *fputc()* function shall write the byte specified by *c* (converted to an **unsigned char**) to the
 13638 output stream pointed to by *stream*, at the position indicated by the associated file-position
 13639 indicator for the stream (if defined), and advances the indicator appropriately. If the file cannot
 13640 support positioning requests, or if the stream was opened with append mode, the byte shall be
 13641 appended to the output stream.

13642 CX The *st_ctime* and *st_mtime* fields of the file shall be marked for update between the successful
 13643 execution of *fputc()* and the next successful completion of a call to *fflush()* or *fclose()* on the same
 13644 stream or a call to *exit()* or *abort()*.

13645 **RETURN VALUE**

13646 Upon successful completion, *fputc()* shall return the value it has written. Otherwise, it shall
 13647 CX return EOF, the error indicator for the stream shall be set, and *errno* shall be set to indicate the
 13648 error.

13649 **ERRORS**

13650 The *fputc()* function shall fail if either the *stream* is unbuffered or the *stream*'s buffer needs to be
 13651 flushed, and:

13652 CX [EAGAIN] The O_NONBLOCK flag is set for the file descriptor underlying *stream* and the
 13653 process would be delayed in the write operation.

13654 CX [EBADF] The file descriptor underlying *stream* is not a valid file descriptor open for
 13655 writing.

13656 CX [EFBIG] An attempt was made to write to a file that exceeds the maximum file size.

13657 XSI [EFBIG] An attempt was made to write to a file that exceeds the process' file size limit.

13658 CX [EFBIG] The file is a regular file and an attempt was made to write at or beyond the
 13659 offset maximum.

13660 CX [EINTR] The write operation was terminated due to the receipt of a signal, and no data
 13661 was transferred.

13662 CX [EIO] A physical I/O error has occurred, or the process is a member of a
 13663 background process group attempting to write to its controlling terminal,
 13664 TOSTOP is set, the process is neither ignoring nor blocking SIGTTOU, and the
 13665 process group of the process is orphaned. This error may also be returned
 13666 under implementation-defined conditions.

13667 CX [ENOSPC] There was no free space remaining on the device containing the file.

13668 CX [EPIPE] An attempt is made to write to a pipe or FIFO that is not open for reading by
 13669 any process. A SIGPIPE signal shall also be sent to the thread.

13670 The *fputc()* function may fail if:

- 13671 CX [ENOMEM] Insufficient storage space is available.
- 13672 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the
13673 capabilities of the device.
- 13674 **EXAMPLES**
- 13675 None.
- 13676 **APPLICATION USAGE**
- 13677 None.
- 13678 **RATIONALE**
- 13679 None.
- 13680 **FUTURE DIRECTIONS**
- 13681 None.
- 13682 **SEE ALSO**
- 13683 *ferror()*, *fopen()*, *getrlimit()*, *putc()*, *puts()*, *setbuf()*, *ulimit()*, the Base Definitions volume of
13684 IEEE Std. 1003.1-200x, <stdio.h>
- 13685 **CHANGE HISTORY**
- 13686 First released in Issue 1. Derived from Issue 1 of the SVID.
- 13687 **Issue 4**
- 13688 In the DESCRIPTION, the text is changed to make it clear that the function writes byte values,
13689 rather than (possibly multi-byte) character values.
- 13690 In the ERRORS section, text is added to indicate that error returns are only generated when
13691 either the stream is unbuffered, or if the stream buffer needs to be flushed.
- 13692 Also in the ERRORS section, in previous issues generation of the [EIO] error depended on
13693 whether or not an implementation supported Job Control. This functionality is now defined as
13694 mandatory.
- 13695 The [ENXIO] error is moved to the list of optional errors, and all the optional errors are marked
13696 as extensions.
- 13697 The description of [EINTR] is amended.
- 13698 The [EFBIG] error is marked to show extensions.
- 13699 **Issue 4, Version 2**
- 13700 In the ERRORS section, the description of [EIO] is updated to include the case where a physical
13701 I/O error occurs.
- 13702 **Issue 5**
- 13703 Large File Summit extensions are added.
- 13704 **Issue 6**
- 13705 Extensions beyond the ISO C standard are now marked.
- 13706 The following new requirements on POSIX implementations derive from alignment with the
13707 Single UNIX Specification:
- 13708
 - The [EIO] and [EFBIG] mandatory error conditions are added.
- 13709
 - The [ENOMEM] and [ENXIO] optional error conditions are added.

13710 NAME

13711 fputs — put a string on a stream

13712 SYNOPSIS

13713 #include <stdio.h>

13714 int fputs(const char *restrict *s*, FILE *restrict *stream*);

13715 DESCRIPTION

13716 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 13717 conflict between the requirements described here and the ISO C standard is unintentional. This
 13718 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

13719 The *fputs()* function shall write the null-terminated string pointed to by *s* to the stream pointed
 13720 to by *stream*. The terminating null byte shall not be written.

13721 CX The *st_ctime* and *st_mtime* fields of the file shall be marked for update between the successful
 13722 execution of *fputs()* and the next successful completion of a call to *fflush()* or *fclose()* on the same
 13723 stream or a call to *exit()* or *abort()*.

13724 RETURN VALUE

13725 Upon successful completion, *fputs()* shall return a non-negative number. Otherwise, it shall
 13726 CX return EOF, set an error indicator for the stream, and set *errno* to indicate the error.

13727 ERRORS

13728 Refer to *fputc()*.

13729 EXAMPLES

13730 **Printing to Standard Output**

13731 The following example gets the current time, converts it to a string using *localtime()* and
 13732 *asctime()*, and prints it to standard output using *fputs()*. It then prints the number of minutes to
 13733 an event for which it is waiting.

```
13734 #include <time.h>
13735 #include <stdio.h>
13736 ...
13737 time_t now;
13738 int minutes_to_event;
13739 ...
13740 time(&now);
13741 printf("The time is ");
13742 fputs(asctime(localtime(&now)), stdout);
13743 printf("There are still %d minutes to the event.\n",
13744       minutes_to_event);
13745 ...
```

13746 APPLICATION USAGE

13747 The *puts()* function appends a <newline> character while *fputs()* does not.

13748 RATIONALE

13749 None.

13750 FUTURE DIRECTIONS

13751 None.

13752 **SEE ALSO**

13753 *fopen()*, *putc()*, *puts()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>

13754 **CHANGE HISTORY**

13755 First released in Issue 1. Derived from Issue 1 of the SVID.

13756 **Issue 4**

13757 In the DESCRIPTION, the words “null character” are replaced by “null byte”, to make it clear
13758 that this function deals solely in byte values.

13759 The following change is incorporated for alignment with the ISO C standard:

- 13760
- The type of argument *s* is changed from **char*** to **const char***.

13761 **Issue 6**

13762 Extensions beyond the ISO C standard are now marked.

13763 The *fputs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

13764 NAME

13765 fputcw — put a wide-character code on a stream

13766 SYNOPSIS

13767 #include <stdio.h>

13768 #include <wchar.h>

13769 wint_t fputcw(wchar_t *wc*, FILE **stream*);

13770 DESCRIPTION

13771 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 13772 conflict between the requirements described here and the ISO C standard is unintentional. This
 13773 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

13774 The *fputcw()* function shall write the character corresponding to the wide-character code *wc* to
 13775 the output stream pointed to by *stream*, at the position indicated by the associated file-position
 13776 indicator for the stream (if defined), and advances the indicator appropriately. If the file cannot
 13777 support positioning requests, or if the stream was opened with append mode, the character is
 13778 appended to the output stream. If an error occurs while writing the character, the shift state of
 13779 the output file is left in an undefined state.

13780 CX The *st_ctime* and *st_mtime* fields of the file shall be marked for update between the successful
 13781 execution of *fputcw()* and the next successful completion of a call to *fflush()* or *fclose()* on the
 13782 same stream or a call to *exit()* or *abort()*.

13783 RETURN VALUE

13784 Upon successful completion, *fputcw()* shall return *wc*. Otherwise, it shall return WEOF, the error
 13785 CX indicator for the stream shall be set set, and *errno* shall be set to indicate the error.

13786 ERRORS

13787 The *fputcw()* function shall fail if either the stream is unbuffered or data in the *stream*'s buffer
 13788 needs to be written, and:

13789 CX [EAGAIN] The O_NONBLOCK flag is set for the file descriptor underlying *stream* and the
 13790 process would be delayed in the write operation.

13791 CX [EBADF] The file descriptor underlying *stream* is not a valid file descriptor open for
 13792 writing.

13793 CX [EFBIG] An attempt was made to write to a file that exceeds the maximum file size or
 13794 the process' file size limit.

13795 CX [EFBIG] The file is a regular file and an attempt was made to write at or beyond the
 13796 offset maximum associated with the corresponding stream.

13797 CX [EINTR] The write operation was terminated due to the receipt of a signal, and no data
 13798 was transferred.

13799 CX [EIO] A physical I/O error has occurred, or the process is a member of a
 13800 background process group attempting to write to its controlling terminal,
 13801 TOSTOP is set, the process is neither ignoring nor blocking SIGTTOU, and the
 13802 process group of the process is orphaned. This error may also be returned
 13803 under implementation-defined conditions.

13804 CX [ENOSPC] There was no free space remaining on the device containing the file.

13805 CX [EPIPE] An attempt is made to write to a pipe or FIFO that is not open for reading by
 13806 any process. A SIGPIPE signal shall also be sent to the thread.

13807 The *fputwc()* function may fail if:

13808 CX [ENOMEM] Insufficient storage space is available.

13809 CX [ENXIO] A request was made of a nonexistent device, or the request was outside the
13810 capabilities of the device.

13811 CX [EILSEQ] The wide-character code *wc* does not correspond to a valid character.

13812 EXAMPLES

13813 None.

13814 APPLICATION USAGE

13815 None.

13816 RATIONALE

13817 None.

13818 FUTURE DIRECTIONS

13819 None.

13820 SEE ALSO

13821 *ferror()*, *fopen()*, *setbuf()*, *ulimit()*, the Base Definitions volume of IEEE Std. 1003.1-200x,
13822 `<stdio.h>`, `<wchar.h>`

13823 CHANGE HISTORY

13824 First released in Issue 4. Derived from the MSE working draft.

13825 Issue 4, Version 2

13826 In the ERRORS section, the description of [EIO] is updated to include the case where a physical
13827 I/O error occurs.

13828 Issue 5

13829 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the type of argument *wc*
13830 is changed from `wint_t` to `wchar_t`.

13831 The Optional Header (OH) marking is removed from `<stdio.h>`.

13832 Large File Summit extensions are added.

13833 Issue 6

13834 Extensions beyond the ISO C standard are now marked.

13835 The following new requirements on POSIX implementations derive from alignment with the
13836 Single UNIX Specification:

- 13837 • The [EFBIG] and [EIO] mandatory error conditions are added.

13838 **NAME**13839 `fputws` — put a wide-character string on a stream13840 **SYNOPSIS**13841 `#include <stdio.h>`13842 `#include <wchar.h>`13843 `int fputws(const wchar_t *restrict ws, FILE *restrict stream);`13844 **DESCRIPTION**

13845 CX The functionality described on this reference page is aligned with the ISO C standard. Any
13846 conflict between the requirements described here and the ISO C standard is unintentional. This
13847 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

13848 The `fputws()` function shall write a character string corresponding to the (null-terminated)
13849 wide-character string pointed to by `ws` to the stream pointed to by `stream`. No character
13850 corresponding to the terminating null wide-character code shall be written.

13851 CX The `st_ctime` and `st_mtime` fields of the file shall be marked for update between the successful
13852 execution of `fputws()` and the next successful completion of a call to `fflush()` or `fclose()` on the
13853 same stream or a call to `exit()` or `abort()`.

13854 **RETURN VALUE**

13855 Upon successful completion, `fputws()` shall return a non-negative number. Otherwise, it shall
13856 CX return `-1`, set an error indicator for the stream, and set `errno` to indicate the error.

13857 **ERRORS**13858 Refer to `fputwc()`.13859 **EXAMPLES**

13860 None.

13861 **APPLICATION USAGE**13862 The `fputws()` function does not append a <newline> character.13863 **RATIONALE**

13864 None.

13865 **FUTURE DIRECTIONS**

13866 None.

13867 **SEE ALSO**13868 `fopen()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdio.h>`, `<wchar.h>`13869 **CHANGE HISTORY**

13870 First released in Issue 4. Derived from the MSE working draft.

13871 **Issue 5**13872 The Optional Header (OH) marking is removed from `<stdio.h>`.13873 **Issue 6**

13874 Extensions beyond the ISO C standard are now marked.

13875 The `fputws()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

13876 **NAME**

13877 fread — binary input

13878 **SYNOPSIS**

13879 #include <stdio.h>

```
13880 size_t fread(void *restrict ptr, size_t size, size_t nitems,
13881 FILE *restrict stream);
```

13882 **DESCRIPTION**

13883 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 13884 conflict between the requirements described here and the ISO C standard is unintentional. This
 13885 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

13886 The *fread()* function shall read into the array pointed to by *ptr* up to *nitems* members whose size
 13887 is specified by *size* in bytes, from the stream pointed to by *stream*. For each object, size calls are
 13888 made to the *fgetc()* function and the results stored, in the order read, in an array of **unsigned**
 13889 **char** exactly overlaying the object. The file position indicator for the stream (if defined) is
 13890 advanced by the number of bytes successfully read. If an error occurs, the resulting value of the
 13891 file position indicator for the stream is indeterminate. If a partial member is read, its value is
 13892 indeterminate.

13893 CX The *fread()* function may mark the *st_atime* field of the file associated with *stream* for update. The
 13894 *st_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,
 13895 *fgetwc()*, *fgetws()*, *fread()*, *fscanf()*, *getc()*, *getchar()*, *gets()*, or *scanf()* using *stream* that returns
 13896 data not supplied by a prior call to *ungetc()* or *ungetwc()*.

13897 **RETURN VALUE**

13898 Upon successful completion, *fread()* shall return the number of members successfully read
 13899 which is less than *nitems* only if a read error or end-of-file is encountered. If *size* or *nitems* is 0,
 13900 *fread()* shall return 0 and the contents of the array and the state of the stream remain unchanged.
 13901 CX Otherwise, if a read error occurs, the error indicator for the stream shall be set, and *errno* shall be
 13902 set to indicate the error.

13903 **ERRORS**13904 Refer to *fgetc()*.13905 **EXAMPLES**13906 **Reading from a Stream**13907 The following example reads a single element from the *fp* stream into the array pointed to by *buf*.

```
13908 #include <stdio.h>
13909 ...
13910 size_t bytes_read;
13911 char buf[100];
13912 FILE *fp;
13913 ...
13914 bytes_read = fread(buf, sizeof(buf), 1, fp);
13915 ...
```

13916 **APPLICATION USAGE**

13917 The *ferror()* or *feof()* functions must be used to distinguish between an error condition and an
 13918 end-of-file condition.

13919 Because of possible differences in member length and byte ordering, files written using *fwrite()*
 13920 are application-dependent, and possibly cannot be read using *fread()* by a different application

13921 or by the same application on a different processor.

13922 **RATIONALE**

13923 None.

13924 **FUTURE DIRECTIONS**

13925 None.

13926 **SEE ALSO**

13927 *feof()*, *ferror()*, *fgetc()*, *fopen()*, *getc()*, *gets()*, *scanf()*, the Base Definitions volume of
13928 IEEE Std. 1003.1-200x, `<stdio.h>`

13929 **CHANGE HISTORY**

13930 First released in Issue 1. Derived from Issue 1 of the SVID.

13931 **Issue 4**

13932 The list of functions that may cause the *st_atime* field to be updated is revised.

13933 The following change is incorporated for alignment with the ISO C standard:

- 13934 • In the RETURN VALUE section, the behavior if *size* or *nitems* is 0 is defined.

13935 **Issue 6**

13936 Extensions beyond the ISO C standard are now marked.

13937 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 13938 • The *fread()* prototype is updated.
- 13939 • The DESCRIPTION is updated to describe how the bytes from a call to *fgetc()* are stored.

13940 **NAME**

13941 free — free allocated memory

13942 **SYNOPSIS**

13943 #include <stdlib.h>

13944 void free(void *ptr);

13945 **DESCRIPTION**

13946 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
13947 conflict between the requirements described here and the ISO C standard is unintentional. This
13948 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

13949 The *free()* function shall cause the space pointed to by *ptr* to be deallocated; that is, made
13950 available for further allocation. If *ptr* is a null pointer, no action shall occur. Otherwise, if the
13951 argument does not match a pointer earlier returned by the *calloc()*, *malloc()*, *posix_memalign()*, or
13952 *realloc()* function, or if the space is deallocated by a call to *free()* or *realloc()*, the behavior is
13953 undefined.

13954 Any use of a pointer that refers to freed space results in undefined behavior.

13955 **RETURN VALUE**13956 The *free()* function shall return no value.13957 **ERRORS**

13958 No errors are defined.

13959 **EXAMPLES**

13960 None.

13961 **APPLICATION USAGE**

13962 There is now no requirement for the implementation to support the inclusion of <malloc.h>.

13963 **RATIONALE**

13964 None.

13965 **FUTURE DIRECTIONS**

13966 None.

13967 **SEE ALSO**13968 *calloc()*, *malloc()*, *realloc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>13969 **CHANGE HISTORY**

13970 First released in Issue 1. Derived from Issue 1 of the SVID.

13971 **Issue 4**

13972 The APPLICATION USAGE section is changed to record that <malloc.h> need no longer be
13973 supported on XSI-conformant systems.

13974 The following change is incorporated for alignment with the ISO C standard:

- 13975 • The DESCRIPTION now states that the behavior is undefined if any use is made of a pointer
13976 that refers to freed space. This was implied but not stated explicitly in Issue 3.

13977 **Issue 4, Version 2**

13978 The DESCRIPTION is updated for X/OPEN UNIX conformance to indicate that the *free()*
13979 function can also be used to free memory allocated by *valloc()*.

13980 **Issue 6**

13981 Reference to the *valloc()* function is removed.

13982 NAME

13983 freeaddrinfo, getaddrinfo — get address information

13984 SYNOPSIS

13985 #include <sys/socket.h>

13986 #include <netdb.h>

13987 void freeaddrinfo(struct addrinfo *ai);

13988 int getaddrinfo(const char *restrict nodename, const char *restrict servname, |

13989 const struct addrinfo *restrict hints, struct addrinfo **restrict res); |

13990 DESCRIPTION

13991 The *freeaddrinfo()* function frees one or more **addrinfo** structures returned by *getaddrinfo()*, along
 13992 with any additional storage associated with those structures. If the *ai_next* field of the structure
 13993 is not null, the entire list of structures is freed. The *freeaddrinfo()* function shall support the
 13994 freeing of arbitrary sublists of an **addrinfo** list originally returned by *getaddrinfo()*.

13995 The *getaddrinfo()* function shall translate the name of a service location (for example, a host
 13996 name) and/or a service name and shall return a set of socket addresses and associated
 13997 information to be used in creating a socket with which to address the specified service.

13998 The *freeaddrinfo()* and *getaddrinfo()* functions shall be thread-safe.

13999 The *nodename* and *servname* arguments are either null pointers or pointers to null-terminated
 14000 strings. One or both of these two arguments shall be supplied by the application as a non-null
 14001 pointer.

14002 The format of a valid name depends on the protocol family or families. If a specific family is not
 14003 given and the name could be interpreted as valid within multiple supported families, the
 14004 implementation shall attempt to resolve the name in all supported families and, in absence of
 14005 errors, one or more results shall be returned.

14006 If the *nodename* argument is not null, it can be a descriptive name or can be an address string. If
 14007 IP6 the specified address family is AF_INET, AF_INET6, or AF_UNSPEC, valid descriptive names
 14008 include host names. If the specified address family is AF_INET or AF_UNSPEC, address strings
 14009 using Internet standard dot notation as specified in *inet_addr()* are valid.

14010 IP6 If the specified address family is AF_INET6 or AF_UNSPEC, standard IPv6 text forms described
 14011 in *inet_ntop()* are valid.

14012 If *nodename* is not null, the requested service location is named by *nodename*; otherwise, the
 14013 requested service location is local to the caller.

14014 If *servname* is null, the call shall return network-level addresses for the specified *nodename*. If
 14015 *servname* is not null, it is a null-terminated character string identifying the requested service. This
 14016 can be either a descriptive name or a numeric representation suitable for use with the address
 14017 IP6 family or families. If the specified address family is AF_INET, AF_INET6, or AF_UNSPEC, the
 14018 service can be specified as a string specifying a decimal port number.

14019 If the *hints* argument is not null, it refers to a structure containing input values that may direct
 14020 the operation by providing options and by limiting the returned information to a specific socket
 14021 type, address family and/or protocol. In this *hints* structure every member other than *ai_flags*,
 14022 *ai_family*, *ai_socktype*, and *ai_protocol* shall be set to zero or a null pointer. A value of
 14023 AF_UNSPEC for *ai_family* means that the caller shall accept any protocol family. A value of zero
 14024 for *ai_socktype* means that the caller shall accept any socket type. A value of zero for *ai_protocol*
 14025 means that the caller shall accept any protocol. If *hints* is a null pointer, the behavior shall be as if
 14026 it referred to a structure containing the value zero for the *ai_flags*, *ai_socktype*, and *ai_protocol*
 14027 fields, and AF_UNSPEC for the *ai_family* field.

14028 **Notes:**

- 14029 1. If the caller handles only TCP and not UDP, for example, then the *ai_protocol*
 14030 member of the *hints* structure should be set to IPPROTO_TCP when
 14031 *getaddrinfo()* is called.
- 14032 2. If the caller handles only IPv4 and not IPv6, then the *ai_family* member of the
 14033 *hints* structure should be set to PF_INET when *getaddrinfo()* is called.

14034 The *ai_flags* field to which the *hints* parameter points shall be set to zero or be the bitwise-
 14035 inclusive OR of one or more of the values AI_PASSIVE, AI_CANONNAME, and
 14036 AI_NUMERICHOST.

14037 If the AI_PASSIVE flag is specified, the returned address information shall be suitable for use in
 14038 binding a socket for accepting incoming connections for the specified service. In this case, if the
 14039 *nodename* argument is null, then the IP address portion of the socket address structure shall be
 14040 set to INADDR_ANY for an IPv4 address or IN6ADDR_ANY_INIT for an IPv6 address. If the
 14041 AI_PASSIVE flag is not specified, the returned address information shall be suitable for a call to
 14042 *connect()* (for a connection-mode protocol) or for a call to *connect()*, *sendto()*, or *sendmsg()* (for a
 14043 connectionless protocol). In this case, if the *nodename* argument is null, then the IP address
 14044 portion of the socket address structure shall be set to the loopback address.

14045 If the AI_CANONNAME flag is specified and the *nodename* argument is not null, the function
 14046 attempts to determine the canonical name corresponding to *nodename* (for example, if *nodename*
 14047 is an alias or shorthand notation for a complete name).

14048 If the AI_NUMERICHOST flag is specified, then a non-null *nodename* string supplied shall be a
 14049 numeric host address string. Otherwise, an [EAI_NONAME] error is returned. This flag prevents
 14050 any type of name resolution service (for example, the DNS) from being invoked.

14051 If the AI_NUMERICSERV flag is specified, then a non-null *servname* string supplied shall be a
 14052 numeric port string. Otherwise, an [EAI_NONAME] error is returned. This flag prevents any
 14053 type of name resolution service (for example, NIS+) from being invoked.

14054 The *ai_socktype* field to which argument *hints* points specifies the socket type for the service, as
 14055 defined in *socket()*. If a specific socket type is not given (for example, a value of zero) and the
 14056 service name could be interpreted as valid with multiple supported socket types, the
 14057 implementation shall attempt to resolve the service name for all supported socket types and, in
 14058 the absence of errors, all possible results shall be returned. A non-zero socket type value shall
 14059 limit the returned information to values with the specified socket type.

14060 If the *ai_family* field to which *hints* points has the value AF_UNSPEC, addresses are returned for
 14061 use with any protocol family that can be used with the specified *nodename* and/or *servname*.
 14062 Otherwise, addresses are returned for use only with the specified protocol family. If *ai_family* is
 14063 not AF_UNSPEC and *ai_protocol* is not zero, then addresses are returned for use only with the
 14064 specified protocol family and protocol; the value of *ai_protocol* is interpreted as in a call to the
 14065 *socket()* function with the corresponding values of *ai_family* and *ai_protocol*.

14066 **RETURN VALUE**

14067 A zero return value for *getaddrinfo()* indicates successful completion; a non-zero return value
 14068 indicates failure. The possible values for the failures are listed in the ERRORS section.

14069 Upon successful return of *getaddrinfo()*, the location to which *res* points refers to a linked list of
 14070 **addrinfo** structures, each of which specifies a socket address and information for use in creating
 14071 a socket with which to use that socket address. The list shall include at least one **addrinfo**
 14072 structure. The *ai_next* field of each structure contains a pointer to the next structure on the list, or
 14073 a null pointer if it is the last structure on the list. Each structure on the list includes values for use
 14074 with a call to the *socket()* function, and a socket address for use with the *connect()* function or, if

14075 the AI_PASSIVE flag was specified, for use with the *bind()* function. The fields *ai_family*,
 14076 *ai_socktype*, and *ai_protocol* are usable as the arguments to the *socket()* function to create a socket
 14077 suitable for use with the returned address. The fields *ai_addr* and *ai_addrlen* are usable as the
 14078 arguments to the *connect()* or *bind()* functions with such a socket, according to the AI_PASSIVE
 14079 flag.

14080 If *nodename* is not null, and if requested by the AI_CANONNAME flag, the *ai_canonname* field of
 14081 the first returned **addrinfo** structure points to a null-terminated string containing the canonical
 14082 name corresponding to the input *nodename*; if the canonical name is not available, then
 14083 *ai_canonname* refers to the *nodename* argument or a string with the same contents. The contents of
 14084 the *ai_flags* field of the returned structures are undefined.

14085 All fields in socket address structures returned by *getaddrinfo()* that are not filled in through an
 14086 explicit argument (for example, *sin6_flowinfo* and *sin_zero*) shall be set to zero.

14087 **Note:** This makes it easier to compare socket address structures.

14088 ERRORS

14089 The *getaddrinfo()* function shall fail and return the corresponding value if:

14090 [EAI_AGAIN] The name could not be resolved at this time. Future attempts may succeed.

14091 [EAI_BADFLAGS]

14092 The *flags* parameter had an invalid value.

14093 [EAI_FAIL] A non-recoverable error occurred when attempting to resolve the name.

14094 [EAI_FAMILY] The address family was not recognized.

14095 [EAI_MEMORY] There was a memory allocation failure when trying to allocate storage for the
 14096 return value.

14097 [EAI_NONAME] The name does not resolve for the supplied parameters.

14098 Neither *nodename* nor *servname* were supplied. At least one of these shall be
 14099 supplied.

14100 [EAI_SERVICE] The service passed was not recognized for the specified socket type.

14101 [EAI_SOCKTYPE]

14102 The intended socket type was not recognized.

14103 [EAI_SYSTEM] A system error occurred; the error code can be found in *errno*.

14104 EXAMPLES

14105 None.

14106 APPLICATION USAGE

14107 None.

14108 RATIONALE

14109 None.

14110 FUTURE DIRECTIONS

14111 None.

14112 SEE ALSO

14113 *connect()*, *gethostbyname()*, *getipnodebyname()*, *getnameinfo()*, *getservbyname()*, *socket()*, the Base
 14114 Definitions volume of IEEE Std. 1003.1-200x, <*netdb.h*>, <*sys/socket.h*>

14115 CHANGE HISTORY

- 14116 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.
- 14117 The **restrict** keyword is added to the *getaddrinfo()* prototype for alignment with the
- 14118 ISO/IEC 9899:1999 standard.

14119 **NAME**

14120 freehostent — network host database functions

14121 **SYNOPSIS**

14122 #include <netdb.h>

14123 void freehostent(struct hostent *ptr);

14124 **DESCRIPTION**14125 Refer to *endhostent()*.

14126 NAME

14127 freopen — open a stream

14128 SYNOPSIS

14129 #include <stdio.h>

14130 FILE *freopen(const char *restrict *filename*, const char *restrict *mode*,
14131 FILE *restrict *stream*);

14132 DESCRIPTION

14133 CX The functionality described on this reference page is aligned with the ISO C standard. Any
14134 conflict between the requirements described here and the ISO C standard is unintentional. This
14135 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.14136 The *freopen()* function shall first attempt to flush the stream and close any file descriptor
14137 associated with *stream*. Failure to flush or close the file successfully shall be ignored. The error
14138 and end-of-file indicators for the stream shall be cleared.14139 The *freopen()* function shall open the file whose path name is the string pointed to by *filename*
14140 and associate the stream pointed to by *stream* with it. The *mode* argument shall be used just as in
14141 *fopen()*.

14142 The original stream shall be closed regardless of whether the subsequent open succeeds.

14143 If *filename* is a null pointer, the *freopen()* function shall attempt to change the mode of the stream
14144 to that specified by *mode*, as if the name of the file currently associated with the stream had been
14145 used. It is implementation-defined which changes of mode are permitted (if any), and under
14146 what circumstances.14147 XSI After a successful call to the *freopen()* function, the orientation of the stream shall be cleared, the
14148 encoding rule shall be cleared, and the associated **mbstate_t** object shall be set to describe an
14149 initial conversion state.14150 CX The largest value that can be represented correctly in an object of type **off_t** shall be established
14151 as the offset maximum in the open file description.

14152 RETURN VALUE

14153 Upon successful completion, *freopen()* shall return the value of *stream*. Otherwise, a null pointer
14154 CX shall be returned, and *errno* shall be set to indicate the error.

14155 ERRORS

14156 The *freopen()* function shall fail if:14157 CX [EACCES] Search permission is denied on a component of the path prefix, or the file
14158 exists and the permissions specified by *mode* are denied, or the file does not
14159 exist and write permission is denied for the parent directory of the file to be
14160 created.14161 CX [EINTR] A signal was caught during *freopen()*.14162 CX [EISDIR] The named file is a directory and *mode* requires write access.14163 CX [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*
14164 argument.

14165 CX [EMFILE] {OPEN_MAX} file descriptors are currently open in the calling process.

14166 CX [ENAMETOOLONG]

14167 The length of the *filename* argument exceeds {PATH_MAX} or a path name
14168 component is longer than {NAME_MAX}.

14169 CX	[ENFILE]	The maximum allowable number of files is currently open in the system.
14170 CX 14171	[ENOENT]	A component of <i>filename</i> does not name an existing file or <i>filename</i> is an empty string.
14172 CX 14173	[ENOSPC]	The directory or file system that would contain the new file cannot be expanded, the file does not exist, and it was to be created.
14174 CX	[ENOTDIR]	A component of the path prefix is not a directory.
14175 CX 14176	[ENXIO]	The named file is a character special or block special file, and the device associated with this special file does not exist.
14177 CX 14178	[EOVERFLOW]	The named file is a regular file and the size of the file cannot be represented correctly in an object of type <code>off_t</code> .
14179 CX 14180	[EROFS]	The named file resides on a read-only file system and <i>mode</i> requires write access.
14181		The <i>freopen()</i> function may fail if:
14182 CX	[EINVAL]	The value of the <i>mode</i> argument is not valid.
14183 CX 14184	[ELOOP]	More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.
14185 CX 14186 14187	[ENAMETOOLONG]	Path name resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.
14188 CX	[ENOMEM]	Insufficient storage space is available.
14189 CX 14190	[ENXIO]	A request was made of a nonexistent device, or the request was outside the capabilities of the device.
14191 CX 14192	[ETXTBSY]	The file is a pure procedure (shared text) file that is being executed and <i>mode</i> requires write access.

14193 EXAMPLES

14194 Directing Standard Output to a File

14195 The following example logs all standard output to the `/tmp/logfile` file.

```
14196 #include <stdio.h>
14197 ...
14198 FILE *fp;
14199 ...
14200 fp = freopen ("/tmp/logfile", "a+", stdout);
14201 ...
```

14202 APPLICATION USAGE

14203 The *freopen()* function is typically used to attach the preopened *streams* associated with *stdin*,
14204 *stdout*, and *stderr* to other files.

14205 RATIONALE

14206 None.

14207 **FUTURE DIRECTIONS**

14208 None.

14209 **SEE ALSO**14210 *fclose()*, *fopen()*, *fdopen()*, *mbsinit()*, the Base Definitions volume of IEEE Std. 1003.1-200x,
14211 `<stdio.h>`14212 **CHANGE HISTORY**

14213 First released in Issue 1. Derived from Issue 1 of the SVID.

14214 **Issue 4**14215 In the DESCRIPTION, the word “name” is replaced by “path name”, to make it clear that the
14216 function is not limited to accepting file names only.

14217 In the ERRORS section:

- 14218 • The description of the [EMFILE] error has been changed to refer to {OPEN_MAX} file
-
- 14219 descriptors rather than {FOPEN_MAX} file descriptors, directories, and message catalogs.
-
- 14220 • The errors [EINVAL], [ENOMEM], and [ETXTBSY] are marked as extensions.
-
- 14221 • The [ENXIO] error is added in the “may fail” section and marked as an extension.

14222 The following change is incorporated for alignment with the ISO C standard:

- 14223 • The type of arguments
- filename*
- and
- mode*
- are changed from
- char***
- to
- const char***
- .

14224 The following change is incorporated for alignment with the FIPS requirements:

- 14225 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path
-
- 14226 name component is larger than {NAME_MAX} is now defined as mandatory and marked as
-
- 14227 an extension.

14228 **Issue 4, Version 2**

14229 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 14230 • It states that [ELOOP] is returned if too many symbolic links are encountered during path
-
- 14231 name resolution.
-
- 14232 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an
-
- 14233 intermediate result of path name resolution of a symbolic link.

14234 **Issue 5**14235 The DESCRIPTION is updated to indicate that the orientation of the stream is cleared and the
14236 conversion state of the stream is set to an initial conversion state by a successful call to the
14237 *freopen()* function.

14238 Large File Summit extensions are added.

14239 **Issue 6**

14240 Extensions beyond the ISO C standard are now marked.

14241 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 14242 • The [ENAMETOOLONG] error is restored as an error dependent on
- `_POSIX_NO_TRUNC`
- .
-
- 14243 This is since behavior may vary from one file system to another.

14244 The following new requirements on POSIX implementations derive from alignment with the
14245 Single UNIX Specification:

- 14246 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file
-
- 14247 description. This change is to support large files.

- 14248 • In the ERRORS section, the [Eoverflow] condition is added. This change is to support
- 14249 large files.
- 14250 • The [ELOOP] mandatory error condition is added.
- 14251 • A second [ENAMETOOLONG] is added as an optional error condition.
- 14252 • The [EINVAL], [ENOMEM], [ENXIO], and [ETXTBSY] optional error conditions are added.
- 14253 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:
- 14254 • The *freopen()* prototype is updated.
- 14255 • The DESCRIPTION is updated.
- 14256 The wording of the mandatory [ELOOP] error condition is updated, and a second optional
- 14257 [ELOOP] error condition is added.

14258 **NAME**

14259 frexp, frexpf, frexpl — extract mantissa and exponent from a double precision number

14260 **SYNOPSIS**

14261 #include <math.h>

14262 double frexp(double *num*, int **exp*);

14263 float frexpf(float *value*, int **exp*);

14264 long double frexpl(long double *value*, int **exp*);

14265 **DESCRIPTION**

14266 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any
 14267 conflict between the requirements described here and the ISO C standard is unintentional. This
 14268 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

14269 These functions breaks a floating-point number into a normalized fraction and an integral power
 14270 of 2. It stores the integer exponent in the **int** object pointed to by *exp*.

14271 An application wishing to check for error situations should set *errno* to 0 before calling *frexp()*. If
 14272 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

14273 **RETURN VALUE**

14274 These functions shall return the value *x*, such that *x* has a magnitude in the interval $[\frac{1}{2}, 1)$ or 0,
 14275 and *num* equals *x* times 2 raised to the power **exp*.

14276 If *num* is 0, both parts of the result shall be 0.

14277 **XSI** If *num* is NaN, NaN shall be returned, *errno* may be set to [EDOM], and the value of **exp* shall be
 14278 unspecified.

14279 If *num* is $\pm\text{Inf}$, *num* shall be returned, *errno* may be set to [EDOM], and the value of **exp* shall be
 14280 unspecified.

14281 **ERRORS**

14282 These functions may fail if:

14283 **XSI** [EDOM] The value of *num* is NaN or $\pm\text{Inf}$.

14284 **XSI** No other errors shall occur.

14285 **EXAMPLES**

14286 None.

14287 **APPLICATION USAGE**

14288 None.

14289 **RATIONALE**

14290 None.

14291 **FUTURE DIRECTIONS**

14292 None.

14293 **SEE ALSO**

14294 *isnan()*, *ldexp()*, *modf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

14295 **CHANGE HISTORY**

14296 First released in Issue 1. Derived from Issue 1 of the SVID.

14297 **Issue 4**

14298 References to *matherr()* are removed.

14299 The name of the first argument is changed from *value* to *num*.

14300 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the
14301 ISO C standard and to rationalize error handling in the mathematics functions.

14302 The return value specified for [EDOM] is marked as an extension.

14303 **Issue 5**

14304 The DESCRIPTION is updated to indicate how an application should check for an error. This
14305 text was previously published in the APPLICATION USAGE section.

14306 **Issue 6**

14307 The *frexpf()* and *frexpl()* functions are added for alignment with the ISO/IEC 9899:1999
14308 standard.

14309 NAME

14310 fscanf, scanf, sscanf — convert formatted input

14311 SYNOPSIS

14312 #include <stdio.h>

14313 int fscanf(FILE *restrict *stream*, const char *restrict *format*, ...);14314 int scanf(const char *restrict *format*, ...);14315 int sscanf(const char *restrict *s*, const char *restrict *format*, ...);

14316 DESCRIPTION

14317 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 14318 conflict between the requirements described here and the ISO C standard is unintentional. This
 14319 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

14320 The *fscanf()* function reads from the named input *stream*. The *scanf()* function reads from the
 14321 standard input stream *stdin*. The *sscanf()* function reads from the string *s*. Each function reads
 14322 bytes, interprets them according to a format, and stores the results in its arguments. Each
 14323 expects, as arguments, a control string *format* described below, and a set of *pointer* arguments
 14324 indicating where the converted input should be stored. The result is undefined if there are
 14325 insufficient arguments for the format. If the format is exhausted while arguments remain, the
 14326 excess arguments are evaluated but are otherwise ignored.

14327 XSI Conversions can be applied to the *n*th argument after the *format* in the argument list, rather than
 14328 to the next unused argument. In this case, the conversion character '*%*' (see below) is replaced
 14329 by the sequence "*%n\$*", where *n* is a decimal integer in the range [1,{NL_ARGMAX}]. This
 14330 feature provides for the definition of format strings that select arguments in an order
 14331 appropriate to specific languages. In format strings containing the "*%n\$*" form of conversion
 14332 specifications, it is unspecified whether numbered arguments in the argument list can be
 14333 referenced from the format string more than once.

14334 The *format* can contain either form of a conversion specification—that is, '*%*' or "*%n\$*"—but the
 14335 two forms cannot normally be mixed within a single *format* string. The only exception to this is
 14336 that "*%%*" or "*%**" can be mixed with the "*%n\$*" form.

14337 The *fscanf()* function in all its forms allows for detection of a language-dependent radix
 14338 character in the input string. The radix character is defined in the program's locale (category
 14339 *LC_NUMERIC*). In the POSIX locale, or in a locale where the radix character is not defined, the
 14340 radix character defaults to a period ('.').

14341 The format is a character string, beginning and ending in its initial shift state, if any, composed
 14342 of zero or more directives. Each directive is composed of one of the following: one or more
 14343 white-space characters (<space>, <tab>, <newline>, <vertical-tab>, or <form-feed> characters);
 14344 an ordinary character (neither '*%*' nor a white-space character); or a conversion specification.
 14345 XSI Each conversion specification is introduced by the character '*%*' or the character sequence
 14346 "*%n\$*", after which the following appear in sequence:

- 14347 • An optional assignment-suppressing character '***'.
- 14348 • An optional non-zero decimal integer that specifies the maximum field width.
- 14349 • An option length modifier that specifies the size of the receiving object.
- 14350 • A conversion character that specifies the type of conversion to be applied. The valid
 14351 conversion characters are described below.

14352 The *fscanf()* functions execute each directive of the format in turn. If a directive fails, as detailed
 14353 below, the function shall return. Failures are described as input failures (due to the
 14354 unavailability of input bytes) or matching failures (due to inappropriate input).

14355 A directive composed of one or more white-space characters is executed by reading input until
 14356 no more valid input can be read, or up to the first byte which is not a white-space character,
 14357 which remains unread.

14358 A directive that is an ordinary character is executed as follows: the next byte is read from the
 14359 input and compared with the byte that comprises the directive; if the comparison shows that
 14360 they are not equivalent, the directive fails, and the differing and subsequent bytes remain
 14361 unread. Similarly, if end-of-file, an encoding error, or a read error prevents a character from
 14362 being read, the directive fails.

14363 A directive that is a conversion specification defines a set of matching input sequences, as
 14364 described below for each conversion character. A conversion specification is executed in the
 14365 following steps:

14366 Input white-space characters (as specified by *isspace()*) are skipped, unless the conversion
 14367 specification includes a ' [', c, C, or n conversion character.

14368 An item is read from the input, unless the conversion specification includes an n conversion
 14369 character. An input item is defined as the longest sequence of input bytes (up to any specified
 14370 maximum field width, which may be measured in characters or bytes dependent on the
 14371 conversion character) which is an initial subsequence of a matching sequence. The first byte, if
 14372 any, after the input item remains unread. If the length of the input item is 0, the execution of the
 14373 conversion specification fails; this condition is a matching failure, unless end-of-file, an encoding
 14374 error, or a read error prevented input from the stream, in which case it is an input failure.

14375 Except in the case of a '%' conversion character, the input item (or, in the case of a %n
 14376 conversion specification, the count of input bytes) is converted to a type appropriate to the
 14377 conversion character. If the input item is not a matching sequence, the execution of the
 14378 conversion specification fails; this condition is a matching failure. Unless assignment
 14379 suppression was indicated by a '*' , the result of the conversion is placed in the object pointed
 14380 to by the first argument following the *format* argument that has not already received a
 14381 XSI conversion result if the conversion specification is introduced by '%', or in the *nth* argument if
 14382 introduced by the character sequence "%n\$". If this object does not have an appropriate type, or
 14383 if the result of the conversion cannot be represented in the space provided, the behavior is
 14384 undefined.

14385 The length modifiers and their meanings are:

14386 *hh* Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument
 14387 with type pointer to **signed char** or **unsigned char**.

14388 *h* Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument
 14389 with type pointer to **short** or **unsigned short**.

14390 *l* (ell) Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument
 14391 with type pointer to **long** or **unsigned long**; that a following *a, A, e, E, f, F, g,* or *G*
 14392 conversion specifier applies to an argument with type pointer to **double**; or that a
 14393 following *c, s,* or ' [' conversion specifier applies to an argument with type pointer to
 14394 **wchar_t**.

14395 *ll* (ell-ell) Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument
 14396 with type pointer to **long long** or **unsigned long long**.

14397 *j* Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument
 14398 with type pointer to **intmax_t** or **uintmax_t**.

14399 *z* Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument
 14400 with type pointer to **size_t** or the corresponding signed integer type.

14401	<i>t</i>	Specifies that a following <i>d</i> , <i>i</i> , <i>o</i> , <i>u</i> , <i>x</i> , <i>X</i> , or <i>n</i> conversion specifier applies to an argument with type pointer to ptrdiff_t or the corresponding unsigned type.
14402		
14403	<i>L</i>	Specifies that a following <i>a</i> , <i>A</i> , <i>e</i> , <i>E</i> , <i>f</i> , <i>F</i> , <i>g</i> , or <i>G</i> conversion specifier applies to an argument with type pointer to long double .
14404		
14405		If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.
14406		
14407		The following conversion characters are valid:
14408	<i>d</i>	Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of <i>strtol()</i> with the value 10 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to int .
14409		
14410		
14411		
14412	<i>i</i>	Matches an optionally signed integer, whose format is the same as expected for the subject sequence of <i>strtol()</i> with 0 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to int .
14413		
14414		
14415		
14416	<i>o</i>	Matches an optionally signed octal integer, whose format is the same as expected for the subject sequence of <i>strtoul()</i> with the value 8 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to unsigned .
14417		
14418		
14419		
14420	<i>u</i>	Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of <i>strtoul()</i> with the value 10 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to unsigned .
14421		
14422		
14423		
14424	<i>x</i>	Matches an optionally signed hexadecimal integer, whose format is the same as expected for the subject sequence of <i>strtoul()</i> with the value 16 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to unsigned .
14425		
14426		
14427		
14428	<i>a</i> , <i>e</i> , <i>f</i> , <i>g</i>	Matches an optionally signed floating-point number, infinity, or NaN, whose format is the same as expected for the subject sequence of <i>strtod()</i> . In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to float .
14429		
14430		
14431		
14432		If the <i>fprintf()</i> family of functions generates character string representations for infinity and NaN (a symbolic entity encoded in floating-point format) to support IEEE Std. 754-1985, the <i>fscanf()</i> family of functions shall recognize them as input.
14433		
14434		
14435	<i>s</i>	Matches a sequence of bytes that are not white-space characters. The application shall ensure that the corresponding argument is a pointer to the initial byte of an array of char , signed char , or unsigned char large enough to accept the sequence and a terminating null character code, which shall be added automatically.
14436		
14437		
14438		
14439		If an <i>l</i> (ell) qualifier is present, the input is a sequence of characters that begins in the initial shift state. Each character is converted to a wide character as if by a call to the <i>mbrtowc()</i> function, with the conversion state described by an mbstate_t object initialized to zero before the first character is converted. The application shall ensure that the corresponding argument is a pointer to an array of wchar_t large enough to accept the sequence and the terminating null wide character, which shall be added automatically.
14440		
14441		
14442		
14443		
14444		
14445		

14446 14447 14448 14449 14450	[Matches a non-empty sequence of bytes from a set of expected bytes (the <i>scanset</i>). The normal skip over white-space characters is suppressed in this case. The application shall ensure that the corresponding argument is a pointer to the initial byte of an array of char , signed char , or unsigned char large enough to accept the sequence and a terminating null byte, which shall be added automatically.
14451 14452 14453 14454 14455 14456 14457		If an <i>l</i> (ell) qualifier is present, the input is a sequence of characters that begins in the initial shift state. Each character in the sequence is converted to a wide character as if by a call to the <i>mbrtowc</i> () function, with the conversion state described by an mbstate_t object initialized to zero before the first character is converted. The application shall ensure that the corresponding argument is a pointer to an array of wchar_t large enough to accept the sequence and the terminating null wide character, which shall be added automatically.
14458 14459 14460 14461 14462 14463 14464 14465 14466 14467 14468		The conversion specification includes all subsequent bytes in the <i>format</i> string up to and including the matching right square bracket (']'). The bytes between the square brackets (the <i>scanlist</i>) comprise the scanset, unless the byte after the left square bracket is a circumflex ('^'), in which case the scanset contains all bytes that do not appear in the scanlist between the circumflex and the right square bracket. If the conversion specification begins with "[]" or "[^]", the right square bracket is included in the scanlist and the next right square bracket is the matching right square bracket that ends the conversion specification; otherwise, the first right square bracket is the one that ends the conversion specification. If a '-' is in the scanlist and is not the first character, nor the second where the first character is a '^', nor the last character, the behavior is implementation-defined.
14469 14470 14471 14472 14473	c	Matches a sequence of bytes of the number specified by the field width (1 if no field width is present in the conversion specification). The application shall ensure that the corresponding argument is a pointer to the initial byte of an array of char , signed char , or unsigned char large enough to accept the sequence. No null byte is added. The normal skip over white-space characters is suppressed in this case.
14474 14475 14476 14477 14478 14479 14480		If an <i>l</i> (ell) qualifier is present, the input is a sequence of characters that begins in the initial shift state. Each character in the sequence is converted to a wide character as if by a call to the <i>mbrtowc</i> () function, with the conversion state described by an mbstate_t object initialized to zero before the first character is converted. The application shall ensure that the corresponding argument is a pointer to an array of wchar_t large enough to accept the resulting sequence of wide characters. No null wide character is added.
14481 14482 14483 14484 14485 14486 14487	p	Matches an implementation-defined set of sequences, which shall be the same as the set of sequences that is produced by the <i>%p</i> conversion of the corresponding <i>fprintf</i> () functions. The application shall ensure that the corresponding argument is a pointer to a pointer to void . The interpretation of the input item is implementation-defined. If the input item is a value converted earlier during the same program execution, the pointer that results shall compare equal to that value; otherwise, the behavior of the <i>%p</i> conversion is undefined.
14488 14489 14490 14491 14492 14493 14494	n	No input is consumed. The application shall ensure that the corresponding argument is a pointer to the integer into which is to be written the number of bytes read from the input so far by this call to the <i>fscanf</i> () functions. Execution of a <i>%n</i> conversion specification does not increment the assignment count returned at the completion of execution of the function. No argument is converted, but one is consumed. If the conversion specification includes an assignment-suppressing character or a field width, the behavior is undefined.

14495 XSI **C** Same as *lc*.

14496 XSI **S** Same as *ls*.

14497 % Matches a single '%'; no conversion or assignment occurs. The complete conversion
14498 specification shall be "%%".

14499 If a conversion specification is invalid, the behavior is undefined.

14500 The conversion characters *A*, *E*, *F*, *G*, and *X* are also valid and behave the same as, respectively, *a*,
14501 *e*, *f*, *g*, and *x*.

14502 If end-of-file is encountered during input, conversion is terminated. If end-of-file occurs before
14503 any bytes matching the current conversion specification (except for %*n*) have been read (other
14504 than leading white-space characters, where permitted), execution of the current conversion
14505 specification terminates with an input failure. Otherwise, unless execution of the current
14506 conversion specification is terminated with a matching failure, execution of the following
14507 conversion specification (if any) is terminated with an input failure.

14508 Reaching the end of the string in *sscanf()* is equivalent to encountering end-of-file for *fscanf()*.

14509 If conversion terminates on a conflicting input, the offending input is left unread in the input.
14510 Any trailing white space (including newline characters) is left unread unless matched by a
14511 conversion specification. The success of literal matches and suppressed assignments is only
14512 directly determinable via the %*n* conversion specification.

14513 The *fscanf()* and *scanf()* functions may mark the *st_atime* field of the file associated with *stream*
14514 for update. The *st_atime* field shall be marked for update by the first successful execution of
14515 *fgetc()*, *fgets()*, *fread()*, *getc()*, *getchar()*, *gets()*, *fscanf()*, or *scanf()* using *stream* that returns data
14516 not supplied by a prior call to *ungetc()*.

14517 **RETURN VALUE**

14518 Upon successful completion, these functions shall return the number of successfully matched
14519 and assigned input items; this number can be 0 in the event of an early matching failure. If the
14520 input ends before the first matching failure or conversion, EOF shall be returned. If a read error
14521 CX occurs, the error indicator for the stream is set, EOF shall be returned, and *errno* shall be set to
14522 indicate the error.

14523 **ERRORS**

14524 For the conditions under which the *fscanf()* functions fail and may fail, refer to *fgetc()* or
14525 *fgetwc()*.

14526 In addition, *fscanf()* may fail if:

14527 XSI **[EILSEQ]** Input byte sequence does not form a valid character.

14528 XSI **[EINVAL]** There are insufficient arguments.

14529 **EXAMPLES**

14530 The call:

14531 `int i, n; float x; char name[50];`
14532 `n = scanf("%d%f%s", &i, &x, name);`

14533 with the input line:

14534 `25 54.32E-1 Hamster`

14535 assigns to *n* the value 3, to *i* the value 25, to *x* the value 5.432, and *name* contains the string
14536 "Hamster".

14537 The call:

```
14538 int i; float x; char name[50];
14539 (void) scanf("%2d%f%d %[0123456789]", &i, &x, name);
```

14540 with input:

```
14541 56789 0123 56a72
```

14542 assigns 56 to *i*, 789.0 to *x*, skips 0123, and places the string "56\0" in *name*. The next call to
14543 *getchar()* shall return the character 'a'.

14544 **Reading Data into an Array**

14545 The following call uses *fscanf()* to read three floating point numbers from standard input into
14546 the *input* array.

```
14547 float input[3];
14548 fscanf (stdin, "%f %f %f", input, input+1, input+2);
```

14549 **APPLICATION USAGE**

14550 If the application calling *fscanf()* has any objects of type **wint_t** or **wchar_t**, it must also include
14551 the **<wchar.h>** header to have these objects defined.

14552 **RATIONALE**

14553 None.

14554 **FUTURE DIRECTIONS**

14555 None.

14556 **SEE ALSO**

14557 *getc()*, *printf()*, *setlocale()*, *strtod()*, *strtol()*, *strtoul()*, *wcrtomb()*, the Base Definitions volume of
14558 IEEE Std. 1003.1-200x, **<langinfo.h>**, **<stdio.h>**, **<wchar.h>**, the Base Definitions volume of
14559 IEEE Std. 1003.1-200x, Chapter 7, Locale

14560 **CHANGE HISTORY**

14561 First released in Issue 1. Derived from Issue 1 of the SVID.

14562 **Issue 4**

14563 Use of the terms “byte” and “character” is rationalized to make it clear when single-byte and
14564 multi-byte values can be used. Similarly, use of the terms “conversion specification” and
14565 “conversion character” is now more precise.

14566 Various errors are corrected. For example, the description of the *d* conversion character
14567 contained an erroneous reference to *strtod()* in Issue 3. This is replaced in this issue by reference
14568 to *strtol()*.

14569 The DESCRIPTION is updated in a number of places to indicate further implications of the
14570 “%n\$” form of a conversion. All references to this functionality, which is not specified in the
14571 ISO C standard, are marked as extensions.

14572 The ERRORS section is changed to refer to the entries for *fgetc()* and *fgetwc()*, the [EINVAL]
14573 error is marked as an extension, and the [EILSEQ] error is added and marked as an extension.

14574 The detailed description of this function including the CHANGE HISTORY section for *scanf()* is
14575 provided here instead of under *scanf()*.

14576 The APPLICATION USAGE section is amended to record the need for **<sys/types.h>** or
14577 **<stddef.h>** if type **wchar_t** is required.

- 14578 The following changes are incorporated for alignment with the ISO C standard:
- 14579 • The type of the argument *format* for all functions, and the type of argument *s* for *sscanf()*, are
14580 changed from **char*** to **const char***.
 - 14581 • The description is updated in various places to align more closely with the text of the ISO C
14582 standard. In particular, this issue fully defines the *L* conversion character, allows for the
14583 support of multi-byte coded character sets (although these are not mandated by X/Open),
14584 and fills in a number of gaps in the definition (for example, by defining termination
14585 conditions for *sscanf()*).
 - 14586 • Following an ANSI interpretation, the effect of conversion specifications that consume no
14587 input is better defined, and is no longer marked as an extension.
- 14588 The following change is incorporated for alignment with the MSE working draft.
- 14589 • The *C* and *S* conversion characters are added, indicating a pointer in the argument list to the
14590 initial wide-character code of an array large enough to accept the input sequence.
- 14591 **Issue 5**
- 14592 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the *l* (ell) qualifier is now
14593 defined for the *c*, *s*, and *' ['* conversion characters.
- 14594 The DESCRIPTION is updated to indicate that if infinity and NaN can be generated by the
14595 *fprintf()* family of functions, then they are recognized by the *fscanf()* family.
- 14596 **Issue 6**
- 14597 The Open Group corrigenda items U021/7 and U028/10 have been applied. These correct
14598 several occurrences of “characters” in the text which have been replaced with the term “bytes”.
- 14599 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 14600 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:
- 14601 • The prototypes for *fscanf()*, *scanf()*, and *sscanf()* are updated.
 - 14602 • The DESCRIPTION is updated.

14603 **NAME**

14604 fseek, fseeko — reposition a file-position indicator in a stream

14605 **SYNOPSIS**

14606 #include <stdio.h>

14607 int fseek(FILE *stream, long offset, int whence);

14608 CX int fseeko(FILE *stream, off_t offset, int whence);

14609

14610 **DESCRIPTION**

14611 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 14612 conflict between the requirements described here and the ISO C standard is unintentional. This
 14613 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

14614 The *fseek()* function shall set the file-position indicator for the stream pointed to by *stream*. If a
 14615 read or write error occurs, the error indicator for the stream shall be set and *fseek()* fails.

14616 The new position, measured in bytes from the beginning of the file, shall be obtained by adding
 14617 *offset* to the position specified by *whence*. The specified point is the beginning of the file for
 14618 {SEEK_SET}, the current value of the file-position indicator for {SEEK_CUR}, or end-of-file for
 14619 {SEEK_END}.

14620 If the stream is to be used with wide-character input/output functions, the application shall
 14621 ensure that *offset* is either 0 or a value returned by an earlier call to *ftell()* on the same stream and
 14622 *whence* is {SEEK_SET}.

14623 A successful call to *fseek()* shall clear the end-of-file indicator for the stream and undo any effects
 14624 of *ungetc()* and *ungetwc()* on the same stream. After an *fseek()* call, the next operation on an
 14625 update stream may be either input or output.

14626 If the most recent operation, other than *ftell()*, on a given stream is *flush()*, the file offset in the
 14627 underlying open file description shall be adjusted to reflect the location specified by *fseek()*.

14628 The *fseek()* function shall allow the file-position indicator to be set beyond the end of existing
 14629 data in the file. If data is later written at this point, subsequent reads of data in the gap shall
 14630 return bytes with the value 0 until data is actually written into the gap.

14631 CX The behavior of *fseek()* on devices which are incapable of seeking is implementation-defined.
 14632 The value of the file offset associated with such a device is undefined.

14633 If the stream is writable and buffered data had not been written to the underlying file, *fseek()*
 14634 shall cause the unwritten data to be written to the file and shall mark the *st_ctime* and *st_mtime*
 14635 fields of the file for update.

14636 In a locale with state-dependent encoding, whether *fseek()* restores the stream's shift state is
 14637 implementation-defined.

14638 The *fseeko()* function is equivalent to the *fseek()* function except that the *offset* argument is of
 14639 type **off_t**.

14640 **RETURN VALUE**14641 CX The *fseek()* and *fseeko()* functions shall return 0 if they succeed.14642 CX Otherwise, they shall return -1 and set *errno* to indicate the error.14643 **ERRORS**

14644 CX The *fseek()* and *fseeko()* functions shall fail if, either the *stream* is unbuffered or the *stream*'s
 14645 CX buffer needed to be flushed, and the call to *fseek()* or *fseeko()* causes an underlying *lseek()* or
 14646 *write()* to be invoked:

14647 CX 14648	[EAGAIN]	The O_NONBLOCK flag is set for the file descriptor and the process would be delayed in the write operation.
14649 CX 14650	[EBADF]	The file descriptor underlying the stream file is not open for writing or the stream's buffer needed to be flushed and the file is not open.
14651 CX	[EFBIG]	An attempt was made to write a file that exceeds the maximum file size.
14652 XSI	[EFBIG]	An attempt was made to write a file that exceeds the process' file size limit.
14653 CX 14654	[EFBIG]	The file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.
14655 CX 14656	[EINTR]	The write operation was terminated due to the receipt of a signal, and no data was transferred.
14657 CX 14658	[EINVAL]	The <i>whence</i> argument is invalid. The resulting file-position indicator would be set to a negative value.
14659 CX 14660 14661 14662 14663	[EIO]	A physical I/O error has occurred, or the process is a member of a background process group attempting to perform a <i>write()</i> to its controlling terminal, TOSTOP is set, the process is neither ignoring nor blocking SIGTTOU, and the process group of the process is orphaned. This error may also be returned under implementation-defined conditions.
14664 CX	[ENOSPC]	There was no free space remaining on the device containing the file.
14665 CX 14666	[EOVERFLOW]	For <i>fseek()</i> , the resulting file offset would be a value which cannot be represented correctly in an object of type long .
14667 CX 14668	[EOVERFLOW]	For <i>fseeko()</i> , the resulting file offset would be a value which cannot be represented correctly in an object of type off_t .
14669 CX	[EPIPE]	The file descriptor underlying <i>stream</i> is associated with a pipe or FIFO.
14670 CX 14671	[EPIPE]	An attempt was made to write to a pipe or FIFO that is not open for reading by any process; a SIGPIPE signal shall also be sent to the thread.
14672 CX 14673	[ENXIO]	A request was made of a nonexistent device, or the request was outside the capabilities of the device.

14674 **EXAMPLES**

14675 None.

14676 **APPLICATION USAGE**

14677 None.

14678 **RATIONALE**

14679 None.

14680 **FUTURE DIRECTIONS**

14681 None.

14682 **SEE ALSO**

14683 *fopen()*, *fsetpos()*, *ftell()*, *getrlimit()*, *rewind()*, *ulimit()*, *ungetc()*, the Base Definitions volume of
 14684 IEEE Std. 1003.1-200x, <stdio.h>

14685 **CHANGE HISTORY**

14686 First released in Issue 1. Derived from Issue 1 of the SVID.

14687 Issue 4

14688 In the DESCRIPTION, the words “The *seek()* function does not, by itself, extend the size of a
14689 file” are deleted.

14690 In the RETURN VALUE section, the value `-1` is marked as an extension. This is because the
14691 ISO POSIX-1 standard only requires that a non-zero value is returned.

14692 In the ERRORS section, text is added to indicate that error returns are only generated when
14693 either the stream is unbuffered, or if the stream buffer needs to be flushed.

14694 The “fail” and “may fail” parts of the ERRORS section are revised for consistency with *lseek()*
14695 and *write()*.

14696 Text associated with the [EIO] error is expanded and the [ENXIO] error is added.

14697 Text is added to explain how *fseek()* is used with wide-character input/output; this is marked as
14698 a WP extension.

14699 The [EFBIG] error is marked to show extensions.

14700 The APPLICATION USAGE section is added.

14701 The following change is incorporated for alignment with the ISO C standard:

- 14702 • The type of argument *offset* is now defined in full as **long** instead of **long**.

14703 The following change is incorporated for alignment with the FIPS requirements:

- 14704 • The [EINTR] error is no longer an indication that the implementation does not report partial
14705 transfers.

14706 Issue 4, Version 2

14707 In the ERRORS section, the description of [EIO] is updated to include the case where a physical
14708 I/O error occurs.

14709 Issue 5

14710 Normative text previously in the APPLICATION USAGE section is moved to the
14711 DESCRIPTION.

14712 Large File Summit extensions are added.

14713 Issue 6

14714 Extensions beyond the ISO C standard are now marked.

14715 The following new requirements on POSIX implementations derive from alignment with the
14716 Single UNIX Specification:

- 14717 • The *fseeko()* function is added.
- 14718 • The [EFBIG], [EOVERFLOW], and [ENXIO] mandatory error conditions are added.

14719 The following change is incorporated for alignment with the FIPS requirements:

- 14720 • The [EINTR] error is no longer an indication that the implementation does not report partial
14721 transfers.

14722 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

14723 The DESCRIPTION is updated to explicitly state that *fseek()* sets the file-position indicator, and
14724 then on error the error indicator is set and *fseek()* fails. This is for alignment with the
14725 ISO/IEC 9899:1999 standard.

14726 **NAME**

14727 fsetpos — set current file position

14728 **SYNOPSIS**

14729 #include <stdio.h>

14730 int fsetpos(FILE *stream, const fpos_t *pos);

14731 **DESCRIPTION**

14732 CX The functionality described on this reference page is aligned with the ISO C standard. Any
14733 conflict between the requirements described here and the ISO C standard is unintentional. This
14734 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

14735 The *fsetpos()* function shall set the file position and state indicators for the stream pointed to by
14736 *stream* according to the value of the object pointed to by *pos*, which the application shall ensure
14737 is a value obtained from an earlier call to *fgetpos()* on the same stream. If a read or write error
14738 occurs, the error indicator for the stream shall be set and *fsetpos()* fails.

14739 A successful call to the *fsetpos()* function shall clear the end-of-file indicator for the stream and
14740 undo any effects of *ungetc()* on the same stream. After an *fsetpos()* call, the next operation on an
14741 update stream may be either input or output.

14742 CX The behavior of *fsetpos()* on devices which are incapable of seeking is implementation-defined.
14743 The value of the file offset associated with such a device is undefined.

14744 **RETURN VALUE**

14745 The *fsetpos()* function shall return 0 if it succeeds; otherwise, it shall return a non-zero value and
14746 set *errno* to indicate the error.

14747 **ERRORS**14748 The *fsetpos()* function may fail if:

14749 CX [EBADF] The file descriptor underlying *stream* is not valid.

14750 CX [ESPIPE] The file descriptor underlying *stream* is associated with a pipe, FIFO, or socket.

14751

14752 **EXAMPLES**

14753 None.

14754 **APPLICATION USAGE**

14755 None.

14756 **RATIONALE**

14757 None.

14758 **FUTURE DIRECTIONS**

14759 None.

14760 **SEE ALSO**

14761 *fopen()*, *ftell()*, *rewind()*, *ungetc()*, the Base Definitions volume of IEEE Std. 1003.1-200x,
14762 <stdio.h>

14763 **CHANGE HISTORY**

14764 First released in Issue 4. Derived from the ISO C standard.

14765 **Issue 6**

14766 Extensions beyond the ISO C standard are now marked.

14767 An additional [ESPIPE] error condition is added for sockets.

14768	The DESCRIPTION is updated to avoid use of the term “must” for application requirements.	
14769	The DESCRIPTION is updated to clarify that the error indicator is set for the stream on a read or	
14770	write error. This is for alignment with the ISO/IEC 9899: 1999 standard.	

14771 **NAME**

14772 fstat — get file status

14773 **SYNOPSIS**

14774 #include <sys/stat.h>

14775 int fstat(int *fildev*, struct stat **buf*);14776 **DESCRIPTION**14777 The *fstat()* function obtains information about an open file associated with the file descriptor
14778 *fildev*, and writes it to the area pointed to by *buf*.14779 SHM If *fildev* references a shared memory object, the implementation need update in the **stat** structure
14780 pointed to by the *buf* argument only the *st_uid*, *st_gid*, *st_size*, and *st_mode* fields, and only the
14781 S_IRUSR, S_IWUSR, S_IRGRP, S_IWGRP, S_IROTH, and S_IWOTH file permission bits need be
14782 valid.14783 TYM If *fildev* references a typed memory object, the implementation need update in the **stat** structure
14784 pointed to by the *buf* argument only the *st_uid*, *st_gid*, *st_size*, and *st_mode* fields, and only the
14785 S_IRUSR, S_IWUSR, S_IRGRP, S_IWGRP, S_IROTH, and S_IWOTH file permission bits need be
14786 valid.14787 The *buf* argument is a pointer to a **stat** structure, as defined in <sys/stat.h>, into which
14788 information is placed concerning the file.14789 The structure members *st_mode*, *st_ino*, *st_dev*, *st_uid*, *st_gid*, *st_atime*, *st_ctime*, and *st_mtime*
14790 shall have meaningful values for all file types defined in this volume of IEEE Std. 1003.1-200x.
14791 The value of the member *st_nlink* shall be set to the number of links to the file.14792 An implementation that provides additional or alternative file access control mechanisms may,
14793 under implementation-defined conditions, cause *fstat()* to fail.14794 The *fstat()* function updates any time-related fields as described in **File Times Update** (see the
14795 Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 6, Character Set), before writing into
14796 the **stat** structure.14797 **RETURN VALUE**14798 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to
14799 indicate the error.14800 **ERRORS**14801 The *fstat()* function shall fail if:14802 [EBADF] The *fildev* argument is not a valid file descriptor.

14803 [EIO] An I/O error occurred while reading from the file system.

14804 [EOVERFLOW] The file size in bytes or the number of blocks allocated to the file or the file
14805 serial number cannot be represented correctly in the structure pointed to by
14806 *buf*.14807 The *fstat()* function may fail if:14808 [EOVERFLOW] One of the values is too large to store into the structure pointed to by the *buf*
14809 argument.

14810 **EXAMPLES**14811 **Obtaining File Status Information**

14812 The following example shows how to obtain file status information for a file named
 14813 `/home/cnd/mod1`. The structure variable `buffer` is defined for the `stat` structure. The
 14814 `/home/cnd/mod1` file is opened with read/write privileges and is passed to the open file
 14815 descriptor `fildev`.

```
14816 #include <sys/types.h>
14817 #include <sys/stat.h>
14818 #include <fcntl.h>

14819 struct stat buffer;
14820 int status;
14821 ...
14822 fildev = open("/home/cnd/mod1", O_RDWR);
14823 status = fstat(fildev, &buffer);
```

14824 **APPLICATION USAGE**

14825 None.

14826 **RATIONALE**

14827 None.

14828 **FUTURE DIRECTIONS**

14829 None.

14830 **SEE ALSO**

14831 `lstat()`, `stat()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<sys/stat.h>`, `<sys/types.h>`

14832 **CHANGE HISTORY**

14833 First released in Issue 1. Derived from Issue 1 of the SVID.

14834 **Issue 4**

14835 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on
 14836 XSI-conformant systems.

14837 The following changes are incorporated in the DESCRIPTION for alignment with the
 14838 ISO POSIX-1 standard:

- 14839 • A paragraph defining the contents of `stat` structure members is added.
- 14840 • The words “extended security controls” are replaced by “additional or alternative file access
 14841 control mechanisms”.

14842 **Issue 4, Version 2**

14843 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 14844 • The [EIO] error is added as a mandatory error indicated the occurrence of an I/O error.
- 14845 • The [EOVERFLOW] error is added as an optional error indicating that one of the values is
 14846 too large to store in the area pointed to by `buf`.

14847 **Issue 5**

14848 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.

14849 Large File Summit extensions are added.

14850 **Issue 6**

14851 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

14852 The following new requirements on POSIX implementations derive from alignment with the
14853 Single UNIX Specification:

14854 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was
14855 required for conforming implementations of previous POSIX specifications, it was not
14856 required for UNIX applications.

14857 • The [EIO] mandatory error condition is added.

14858 • The [EOVERFLOW] mandatory error condition is added. This change is to support large
14859 files.

14860 • The [EOVERFLOW] optional error condition is added.

14861 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that
14862 shared memory object semantics apply to typed memory objects.

14863 **NAME**

14864 fstatvfs, statvfs — get file system information

14865 **SYNOPSIS**

14866 XSI #include <sys/statvfs.h>

14867 int fstatvfs(int *fildev*, struct statvfs **buf*);14868 int statvfs(const char *restrict *path*, struct statvfs *restrict *buf*);

14869

14870 **DESCRIPTION**14871 The *fstatvfs()* function obtains information about the file system containing the file referenced by
14872 *fildev*.14873 The following flags can be returned in the *f_flag* member:

14874 ST_RDONLY Read-only file system.

14875 ST_NOSUID Setuid/setgid bits ignored by *exec*.14876 The *statvfs()* function obtains descriptive information about the file system containing the file
14877 named by *path*.14878 For both functions, the *buf* argument is a pointer to a **statvfs** structure that shall be filled. Read,
14879 write, or execute permission of the named file is not required.14880 It is unspecified whether all members of the **statvfs** structure have meaningful values on all file
14881 systems.14882 **RETURN VALUE**14883 Upon successful completion, *statvfs()* shall return 0. Otherwise, it shall return -1 and set *errno* to
14884 indicate the error.14885 **ERRORS**14886 The *fstatvfs()* and *statvfs()* functions shall fail if:

14887 [EIO] An I/O error occurred while reading the file system.

14888 [EINTR] A signal was caught during execution of the function.

14889 [EOVERFLOW] One of the values to be returned cannot be represented correctly in the
14890 structure pointed to by *buf*.14891 The *fstatvfs()* function shall fail if:14892 [EBADF] The *fildev* argument is not an open file descriptor.14893 The *statvfs()* function shall fail if:

14894 [EACCES] Search permission is denied on a component of the path prefix.

14895 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*
14896 argument.

14897 [ENAMETOOLONG]

14898 The length of a path name exceeds {PATH_MAX} or a path name component
14899 is longer than {NAME_MAX}.14900 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.14901 [ENOTDIR] A component of the path prefix of *path* is not a directory.14902 The *statvfs()* function may fail if:

14903 [ELOOP] More than {SYMLOOP_MAX} symbolic links were encountered during
 14904 resolution of the *path* argument.

14905 [ENAMETOOLONG]
 14906 Path name resolution of a symbolic link produced an intermediate result
 14907 whose length exceeds {PATH_MAX}.

14908 EXAMPLES

14909 Obtaining File System Information Using *fstatvfs()*

14910 The following example shows how to obtain file system information for the file system upon
 14911 which the file named */home/cnd/mod1* resides, using the *fstatvfs()* function. The
 14912 */home/cnd/mod1* file is opened with read/write privileges and the open file descriptor is passed
 14913 to the *fstatvfs()* function.

```
14914 #include <statvfs.h>
14915 #include <fcntl.h>
14916 struct statvfs buffer;
14917 int status;
14918 ...
14919 fildes = open("/home/cnd/mod1", O_RDWR);
14920 status = fstatvfs(fildes, &buffer);
```

14921 Obtaining File System Information Using *statvfs()*

14922 The following example shows how to obtain file system information for the file system upon
 14923 which the file named */home/cnd/mod1* resides, using the *statvfs()* function.

```
14924 #include <statvfs.h>
14925 struct statvfs buffer;
14926 int status;
14927 ...
14928 status = statvfs("/home/cnd/mod1", &buffer);
```

14929 APPLICATION USAGE

14930 None.

14931 RATIONALE

14932 None.

14933 FUTURE DIRECTIONS

14934 None.

14935 SEE ALSO

14936 *chmod()*, *chown()*, *creat()*, *dup()*, *exec*, *fcntl()*, *link()*, *mknod()*, *open()*, *pipe()*, *read()*, *time()*,
 14937 *unlink()*, *utime()*, *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/statvfs.h>

14938 CHANGE HISTORY

14939 First released in Issue 4, Version 2.

14940 Issue 5

14941 Moved from X/OPEN UNIX extension to BASE.

14942 Large File Summit extensions are added.

14943 **Issue 6**

- 14944 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 14945 The **restrict** keyword is added to the *statvfs()* prototype for alignment with the
14946 ISO/IEC 9899:1999 standard.
- 14947 The wording of the mandatory [ELOOP] error condition is updated, and a second optional
14948 [ELOOP] error condition is added.

14949 **NAME**

14950 fsync — synchronize changes to a file

14951 **SYNOPSIS**

14952 FSC #include <unistd.h>

14953 int fsync(int *fildev*);

14954

14955 **DESCRIPTION**

14956 The *fsync()* function can be used by an application to indicate that all data for the open file
14957 description named by *fildev* is to be transferred to the storage device associated with the file
14958 described by *fildev* in an implementation-defined manner. The *fsync()* function shall not return
14959 until the system has completed that action or until an error is detected.

14960 SIO If `_POSIX_SYNCHRONIZED_IO` is defined, the *fsync()* function shall force all currently queued
14961 I/O operations associated with the file indicated by file descriptor *fildev* to the synchronized I/O
14962 completion state. All I/O operations shall be completed as defined for synchronized I/O file
14963 integrity completion.

14964 **RETURN VALUE**

14965 Upon successful completion, *fsync()* shall return 0. Otherwise, `-1` shall be returned and *errno* set
14966 to indicate the error. If the *fsync()* function fails, outstanding I/O operations are not guaranteed
14967 to have been completed.

14968 **ERRORS**14969 The *fsync()* function shall fail if:14970 [EBADF] The *fildev* argument is not a valid descriptor.14971 [EINTR] The *fsync()* function was interrupted by a signal.14972 [EINVAL] The *fildev* argument does not refer to a file on which this operation is possible.

14973 [EIO] An I/O error occurred while reading from or writing to the file system.

14974 In the event that any of the queued I/O operations fail, *fsync()* shall return the error conditions
14975 defined for *read()* and *write()*.

14976 **EXAMPLES**

14977 None.

14978 **APPLICATION USAGE**

14979 The *fsync()* function should be used by programs which require modifications to a file to be
14980 completed before continuing; for example, a program which contains a simple transaction
14981 facility might use it to ensure that all modifications to a file or files caused by a transaction are
14982 recorded.

14983 **RATIONALE**

14984 The *fsync()* function is intended to force a physical write of data from the buffer cache, and to
14985 assure that after a system crash or other failure that all data up to the time of the *fsync()* call is
14986 recorded on the disk. Since the concepts of “buffer cache”, “system crash”, “physical write”, and
14987 “non-volatile storage” are not defined here, the wording has to be more abstract.

14988 If `_POSIX_SYNCHRONIZED_IO` is not defined, the wording relies heavily on the conformance
14989 document to tell the user what can be expected from the system. It is explicitly intended that a
14990 null implementation is permitted. This could be valid in the case where the system cannot assure
14991 non-volatile storage under any circumstances or when the system is highly fault-tolerant and the
14992 functionality is not required. In the middle ground between these extremes, *fsync()* might or
14993 might not actually cause data to be written where it is safe from a power failure. The

14994 conformance document should identify at least that one configuration exists (and how to obtain
14995 that configuration) where this can be assured for at least some files that the user can select to use
14996 for critical data. It is not intended that an exhaustive list is required, but rather sufficient
14997 information is provided to let the user determine that if he or she has critical data he or she can
14998 configure her system to allow it to be written to non-volatile storage.

14999 It is reasonable to assert that the key aspects of *fsync()* are unreasonable to test in a test suite.
15000 That does not make the function any less valuable, just more difficult to test. A formal
15001 conformance test should probably force a system crash (power shutdown) during the test for
15002 this condition, but it needs to be done in such a way that automated testing does not require this
15003 to be done except when a formal record of the results is being made. It would also not be
15004 unreasonable to omit testing for *fsync()*, allowing it to be treated as a quality-of-implementation
15005 issue.

15006 **FUTURE DIRECTIONS**

15007 None.

15008 **SEE ALSO**

15009 *sync()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

15010 **CHANGE HISTORY**

15011 First released in Issue 3.

15012 **Issue 4**

15013 The <**unistd.h**> header is added to the SYNOPSIS section.

15014 In the APPLICATION USAGE section, the words “require a file to be in a known state” are
15015 replaced by “require modifications to a file to be completed before continuing”.

15016 **Issue 5**

15017 Aligned with *fsync()* in the POSIX Realtime Extension. Specifically, the DESCRIPTION and
15018 RETURN VALUE sections are much expanded, and the ERRORS section is updated to indicate
15019 that *fsync()* can return the error conditions defined for *read()* and *write()*.

15020 **Issue 6**

15021 This function is marked as part of the File Synchronization option.

15022 The following new requirements on POSIX implementations derive from alignment with the
15023 Single UNIX Specification:

- 15024 • The [EINVAL] and [EIO] mandatory error conditions are added.

15025 **NAME**

15026 ftell, ftello — return a file offset in a stream

15027 **SYNOPSIS**

15028 #include <stdio.h>

15029 long ftell(FILE *stream);

15030 CX off_t ftello(FILE *stream);

15031

15032 **DESCRIPTION**

15033 CX The functionality described on this reference page is aligned with the ISO C standard. Any
15034 conflict between the requirements described here and the ISO C standard is unintentional. This
15035 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

15036 The *ftell()* function shall obtain the current value of the file-position indicator for the stream
15037 pointed to by *stream*.

15038 CX The *ftello()* function is identical to *ftell()* except that the return value is of type **off_t**.

15039 **RETURN VALUE**

15040 CX Upon successful completion, *ftell()* and *ftello()* shall return the current value of the file-position
15041 indicator for the stream measured in bytes from the beginning of the file.

15042 CX Otherwise, *ftell()* and *ftello()* shall return -1 , cast to **long** and **off_t** respectively, and set *errno* to
15043 indicate the error.

15044 **ERRORS**

15045 CX The *ftell()* and *ftello()* functions shall fail if:

15046 CX [EBADF] The file descriptor underlying *stream* is not an open file descriptor.

15047 CX [EOVERFLOW] For *ftell()*, the current file offset cannot be represented correctly in an object of
15048 type **long**.

15049 CX [EOVERFLOW] For *ftello()*, the current file offset cannot be represented correctly in an object
15050 of type **off_t**.

15051 CX [ESPIPE] The file descriptor underlying *stream* is associated with a pipe or FIFO.

15052 The *ftell()* function may fail if:

15053 CX [ESPIPE] The file descriptor underlying *stream* is associated with a socket.

15054 **EXAMPLES**

15055 None.

15056 **APPLICATION USAGE**

15057 None.

15058 **RATIONALE**

15059 None.

15060 **FUTURE DIRECTIONS**

15061 None.

15062 **SEE ALSO**15063 *fgetpos()*, *fopen()*, *fseek()*, *lseek()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdio.h**>

15064 **CHANGE HISTORY**

15065 First released in Issue 1. Derived from Issue 1 of the SVID.

15066 **Issue 4**

15067 The following change is incorporated for alignment with the ISO C standard:

- 15068 • The function return value is now defined in full as **long**. It was previously defined as **long**.

15069 **Issue 5**

15070 Large File Summit extensions are added.

15071 **Issue 6**

15072 Extensions beyond the ISO C standard are now marked.

15073 The following new requirements on POSIX implementations derive from alignment with the
15074 Single UNIX Specification:

- 15075 • The *ftello()* function is added.
- 15076 • The [Eoverflow] error conditions are added.

15077 An additional [ESPIPE] error condition is added for sockets.

15078 **NAME**15079 ftime — get date and time (**LEGACY**)15080 **SYNOPSIS**

15081 xSI #include <sys/timeb.h>

15082 int ftime(struct timeb *tp);

15083

15084 **DESCRIPTION**

15085 The *ftime()* function shall set the *time* and *millitm* members of the **timeb** structure pointed to by
 15086 *tp* to contain the seconds and milliseconds portions, respectively, of the current time in seconds
 15087 since the Epoch. The contents of the *timezone* and *dstflag* members of *tp* after a call to *ftime()* are
 15088 unspecified.

15089 The system clock need not have millisecond granularity. Depending on any granularity
 15090 (particularly a granularity of one) renders code non-portable.

15091 **RETURN VALUE**15092 Upon successful completion, the *ftime()* function shall return 0; otherwise, -1 shall be returned.15093 **ERRORS**

15094 No errors are defined.

15095 **EXAMPLES**15096 **Getting the Current Time and Date**

15097 The following example shows how to get the current system time values using the *ftime()*
 15098 function. The **timeb** structure pointed to by *tp* is filled with the current system time values for
 15099 *time* and *millitm*.

15100 #include <sys/timeb.h>

15101 struct timeb tp;

15102 int status;

15103 ...

15104 status = ftime(&tp);

15105 **APPLICATION USAGE**

15106 For applications portability, the *time()* function should be used to determine the current time
 15107 instead of *ftime()*.

15108 **RATIONALE**

15109 None.

15110 **FUTURE DIRECTIONS**

15111 This function may be withdrawn in a future version.

15112 **SEE ALSO**

15113 *ctime()*, *gettimeofday()*, *time()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |
 15114 <sys/timeb.h>

15115 **CHANGE HISTORY**

15116 First released in Issue 4, Version 2.

15117 **Issue 5**

15118 Moved from X/OPEN UNIX extension to BASE.

15119 Normative text previously in the APPLICATION USAGE section is moved to the
 15120 DESCRIPTION.

15121 **Issue 6**

15122 This function is marked LEGACY.

15123 The DESCRIPTION is updated to refer to “seconds since the Epoch” rather than “seconds since
15124 00:00:00 UTC (Coordinated Universal Time), January 1 1970” for consistency with other *time* |
15125 functions.

15126 NAME

15127 ftok — generate an IPC key

15128 SYNOPSIS

15129 xSI #include <sys/ipc.h>

15130 key_t ftok(const char *path, int id);

15131

15132 DESCRIPTION

15133 The *ftok()* function shall return a key based on *path* and *id* that is usable in subsequent calls to
 15134 *msgget()*, *semget()*, and *shmget()*. The application shall ensure that the *path* argument is the path
 15135 name of an existing file that the process is able to *stat()*.

15136 The *ftok()* function shall return the same key value for all paths that name the same file, when
 15137 called with the same *id* value, and return different key values when called with different *id*
 15138 values or with paths that name different files existing on the same file system at the same time. It
 15139 is unspecified whether *ftok()* shall return the same key value when called again after the file
 15140 named by *path* is removed and recreated with the same name.

15141 Only the low order 8-bits of *id* are significant. The behavior of *ftok()* is unspecified if these bits
 15142 are 0.

15143 RETURN VALUE

15144 Upon successful completion, *ftok()* shall return a key. Otherwise, *ftok()* shall return (**key_t**)-1
 15145 and set *errno* to indicate the error.

15146 ERRORS

15147 The *ftok()* function shall fail if:

15148 [EACCES] Search permission is denied for a component of the path prefix. |

15149 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path* |
 15150 argument. |

15151 [ENAMETOOLONG] |
 15152 The length of the *path* argument exceeds {PATH_MAX} or a path name |
 15153 component is longer than {NAME_MAX}. |

15154 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string. |

15155 [ENOTDIR] A component of the path prefix is not a directory. |

15156 The *ftok()* function may fail if:

15157 [ELOOP] More than {SYMLOOP_MAX} symbolic links were encountered during |
 15158 resolution of the *path* argument. |

15159 [ENAMETOOLONG] |
 15160 Path name resolution of a symbolic link produced an intermediate result |
 15161 whose length exceeds {PATH_MAX}. |

15162 **EXAMPLES**15163 **Getting an IPC Key**

15164 The following example gets a unique key that can be used by the IPC functions *semget()*,
 15165 *msgget()*, and *shmget()*. The key returned by *ftok()* for this example is based on the ID value *S*
 15166 and the path name */tmp*.

```
15167 #include <sys/ipc.h>
15168 ...
15169 key_t key;
15170 char *path = "/tmp";
15171 int id = 'S';
15172 key = ftok(path, id);
```

15173 **Saving an IPC Key**

15174 The following example gets a unique key based on the path name */tmp* and the ID value *a*. It
 15175 also assigns the value of the resulting key to the *semkey* variable so that it will be available to a
 15176 later call to *semget()*, *msgget()*, or *shmget()*.

```
15177 #include <sys/ipc.h>
15178 ...
15179 key_t semkey;
15180 if ((semkey = ftok("/tmp", 'a')) == (key_t) -1) {
15181     perror("IPC error: ftok"); exit(1);
15182 }
```

15183 **APPLICATION USAGE**

15184 For maximum portability, *id* should be a single-byte character.

15185 **RATIONALE**

15186 None.

15187 **FUTURE DIRECTIONS**

15188 None.

15189 **SEE ALSO**

15190 *msgget()*, *semget()*, *shmget()*, the Base Definitions volume of IEEE Std. 1003.1-200x, *<sys/ipc.h>*

15191 **CHANGE HISTORY**

15192 First released in Issue 4, Version 2.

15193 **Issue 5**

15194 Moved from X/OPEN UNIX extension to BASE.

15195 **Issue 6**

15196 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

15197 The wording of the mandatory [ELOOP] error condition is updated, and a second optional
 15198 [ELOOP] error condition is added.

15199 **NAME**

15200 ftruncate — truncate a file to a specified length

15201 **SYNOPSIS**

15202 #include <unistd.h>

15203 int ftruncate(int *fildev*, off_t *length*);15204 **DESCRIPTION**15205 If *fildev* is not a valid file descriptor open for writing, the *ftruncate()* function shall fail.

15206 If *fildev* refers to a regular file, the *ftruncate()* function shall cause the size of the file to be truncated to *length*. If the size of the file previously exceeded *length*, the extra data shall no longer be available to reads on the file. If the file previously was smaller than this size, *ftruncate()* shall either increase the size of the file or fail. XSI-conformant systems shall increase the size of the file. If the file size is increased, the extended area shall appear as if it were zero-filled. The value of the seek pointer shall not be modified by a call to *ftruncate()*.

15212 Upon successful completion, if *fildev* refers to a regular file, the *ftruncate()* function shall mark for update the *st_ctime* and *st_mtime* fields of the file and the S_ISUID and S_ISGID bits of the file mode may be cleared. If the *ftruncate()* function is unsuccessful, the file is unaffected.

15215 XSI If the request would cause the file size to exceed the soft file size limit for the process, the request shall fail and the implementation shall generate the SIGXFSZ signal for the process.

15217 If *fildev* refers to a directory, *ftruncate()* shall fail.15218 If *fildev* refers to any other file type, except a shared memory object, the result is unspecified.

15219 SHM If *fildev* refers to a shared memory object, *ftruncate()* shall set the size of the shared memory object to *length*.

15221 MF|SHM If the effect of *ftruncate()* is to decrease the size of a shared memory object or memory mapped file and whole pages beyond the new end were previously mapped, then the whole pages beyond the new end shall be discarded.

15224 MPR If the Memory Protection option is supported, references to discarded pages shall result in the generation of a SIGBUS signal; otherwise, the result of such references is undefined.

15226 MF|SHM If the effect of *ftruncate()* is to increase the size of a shared memory object, it is unspecified if the contents of any mapped pages between the old end-of-file and the new are flushed to the underlying object.

15229 **RETURN VALUE**

15230 Upon successful completion, *ftruncate()* shall return 0; otherwise, -1 shall be returned and *errno* set to indicate the error.

15232 **ERRORS**15233 The *ftruncate()* function shall fail if:

15234 [EINTR] A signal was caught during execution.

15235 [EINVAL] The *length* argument was less than 0.

15236 [EFBIG] or [EINVAL]

15237 The *length* argument was greater than the maximum file size.

15238 XSI [EFBIG] The file is a regular file and *length* is greater than the offset maximum established in the open file description associated with *fildev*.

15240 [EIO] An I/O error occurred while reading from or writing to a file system.

- 15241 [EBADF] or [EINVAL] |
 15242 The *fildest* argument is not a file descriptor open for writing. |
 15243 [EINVAL] The *fildest* argument references a file that was opened without write |
 15244 permission. |
 15245 [EROFS] The named file resides on a read-only file system. |

15246 **EXAMPLES**

15247 None.

15248 **APPLICATION USAGE**

15249 None.

15250 **RATIONALE**

15251 The *ftruncate()* function is part of IEEE Std. 1003.1-200x as it was deemed to be more useful than |
 15252 *truncate()*. The *truncate()* function is provided as an XSI extension. |

15253 **FUTURE DIRECTIONS**

15254 None.

15255 **SEE ALSO**15256 *open()*, *truncate()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**> |15257 **CHANGE HISTORY**

15258 First released in Issue 4, Version 2.

15259 **Issue 5**

15260 Moved from X/OPEN UNIX extension to BASE and aligned with *ftruncate()* in the POSIX |
 15261 Realtime Extension. Specifically, the DESCRIPTION is extensively reworded and [EROFS] is |
 15262 added to the list of mandatory errors that can be returned by *ftruncate()*.

15263 Large File Summit extensions are added.

15264 **Issue 6**15265 The *truncate()* function has been split out into a separate reference page.

15266 The following new requirements on POSIX implementations derive from alignment with the |
 15267 Single UNIX Specification:

- 15268 • The DESCRIPTION is change to indicate that if the file size is changed, and if the file is a |
 15269 regular file, the S_ISUID and S_ISGID bits in the file mode may be cleared.

15270 The following changes were made to align with the IEEE P1003.1a draft standard:

- 15271 • The DESCRIPTION text is updated.

15272 XSI-conformant systems are required to increase the size of the file if the file was previously |
 15273 smaller than the size requested. |

15274 **NAME**

15275 ftrylockfile — stdio locking functions

15276 **SYNOPSIS**

15277 TSF #include <stdio.h>

15278 int ftrylockfile(FILE *file);

15279

15280 **DESCRIPTION**

15281 Refer to *flockfile()*.

15282 **NAME**

15283 ftw — traverse (walk) a file tree

15284 **SYNOPSIS**15285 XSI

```
#include <ftw.h>
```

```
15286 int ftw(const char *path, int (*fn)(const char *,
15287     const struct stat *ptr, int flag), int ndirs);
15288
```

15289 **DESCRIPTION**

15290 The *ftw()* function recursively descends the directory hierarchy rooted in *path*. For each object in
 15291 the hierarchy, *ftw()* shall call the function pointed to by *fn*, passing it a pointer to a null-
 15292 terminated character string containing the name of the object, a pointer to a **stat** structure
 15293 containing information about the object, and an integer. Possible values of the integer, defined
 15294 in the <ftw.h> header, are:

15295 FTW_D For a directory.

15296 FTW_DNR For a directory that cannot be read.

15297 FTW_F For a file.

15298 FTW_SL For a symbolic link (but see also FTW_NS below).

15299 FTW_NS For an object other than a symbolic link on which *stat()* could not successfully be
 15300 executed. If the object is a symbolic link and *stat()* failed, it is unspecified whether
 15301 *ftw()* passes FTW_SL or FTW_NS to the user-supplied function.

15302 If the integer is FTW_DNR, descendants of that directory shall not be processed. If the integer is
 15303 FTW_NS, the **stat** structure contains undefined values. An example of an object that would
 15304 cause FTW_NS to be passed to the function pointed to by *fn* would be a file in a directory with
 15305 read but without execute (search) permission.

15306 The *ftw()* function shall visit a directory before visiting any of its descendants.15307 The *ftw()* function shall use at most one file descriptor for each level in the tree.15308 The argument *ndirs* should be in the range of 1 to {OPEN_MAX}.

15309 The tree traversal shall continue until the tree is exhausted, an invocation of *fn* returns a non-
 15310 zero value, or some error, other than [EACCES], is detected within *ftw()*.

15311 The *ndirs* argument shall specify the maximum number of directory streams or file descriptors
 15312 or both available for use by *ftw()* while traversing the tree. When *ftw()* returns it shall close any
 15313 directory streams and file descriptors it uses not counting any opened by the application-
 15314 supplied *fn* function.

15315 **RETURN VALUE**

15316 If the tree is exhausted, *ftw()* shall return 0. If the function pointed to by *fn* returns a non-zero
 15317 value, *ftw()* shall stop its tree traversal and return whatever value was returned by the function
 15318 pointed to by *fn*. If *ftw()* detects an error, it shall return -1 and set *errno* to indicate the error.

15319 If *ftw()* encounters an error other than [EACCES] (see FTW_DNR and FTW_NS above), it shall
 15320 return -1 and set *errno* to indicate the error. The external variable *errno* may contain any error
 15321 value that is possible when a directory is opened or when one of the *stat* functions is executed on
 15322 a directory or file.

15323 **ERRORS**15324 The *ftw()* function shall fail if:15325 [EACCES] Search permission is denied for any component of *path* or read permission is
15326 denied for *path*.15327 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*
15328 argument.

15329 [ENAMETOOLONG]

15330 The length of the *path* argument exceeds {PATH_MAX} or a path name
15331 component is longer than {NAME_MAX}.15332 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.15333 [ENOTDIR] A component of *path* is not a directory.15334 The *ftw()* function may fail if:15335 [EINVAL] The value of the *ndirs* argument is invalid.15336 [ELOOP] More than {SYMLOOP_MAX} symbolic links were encountered during
15337 resolution of the *path* argument.

15338 [ENAMETOOLONG]

15339 Path name resolution of a symbolic link produced an intermediate result
15340 whose length exceeds {PATH_MAX}.15341 In addition, if the function pointed to by *fn* encounters system errors, *errno* may be set
15342 accordingly.15343 **EXAMPLES**15344 **Walking a Directory Structure**15345 The following example walks the current directory structure, calling the *fn* function for every
15346 directory entry, using at most 10 file descriptors:

```

15347 #include <ftw.h>
15348 ...
15349 if (ftw(".", fn, 10) != 0) {
15350     perror("ftw"); exit(2);
15351 }

```

15352 **APPLICATION USAGE**

15353 The *ftw()* function may allocate dynamic storage during its operation. If *ftw()* is forcibly
 15354 terminated, such as by *longjmp()* or *siglongjmp()* being executed by the function pointed to by *fn*
 15355 or an interrupt routine, *ftw()* does not have a chance to free that storage, so it remains
 15356 permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has
 15357 occurred, and arrange to have the function pointed to by *fn* return a non-zero value at its next
 15358 invocation.

15359 **RATIONALE**

15360 None.

15361 **FUTURE DIRECTIONS**

15362 None.

15363 **SEE ALSO**

15364 *longjmp()*, *lstat()*, *malloc()*, *nftw()*, *opendir()*, *siglongjmp()*, *stat()*, the Base Definitions volume of
 15365 IEEE Std. 1003.1-200x, <ftw.h>, <sys/stat.h>

15366 **CHANGE HISTORY**

15367 First released in Issue 1. Derived from Issue 1 of the SVID.

15368 **Issue 4**

15369 The type of argument *path* is changed from **char*** to **const char***. The argument list for *fn* has
 15370 also been defined.

15371 In the DESCRIPTION, the words “other than [EACCES]” are added to the paragraph describing
 15372 termination conditions for tree traversal.

15373 The following change is incorporated for alignment with the FIPS requirements:

- 15374 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path
 15375 name component is larger than {NAME_MAX} is now defined as mandatory and marked as
 15376 an extension.

15377 **Issue 4, Version 2**

15378 The following changes are incorporated for X/OPEN UNIX conformance:

- 15379 • The DESCRIPTION is updated to describe the use of the FTW_SL and FTW_NS values for a
 15380 symbolic link.
- 15381 • The DESCRIPTION states that *ftw()* uses at most one file descriptor for each level in the tree.
- 15382 • The DESCRIPTION constrains *ndirs* to the range from 1 to {OPEN_MAX}.
- 15383 • The RETURN VALUE section is updated to describe the case where *ftw()* encounters an error
 15384 other than [EACCES].
- 15385 • In the ERRORS section, a second [ENAMETOOLONG] condition is defined that may report
 15386 excessive length of an intermediate result of path name resolution of a symbolic link.

15387 **Issue 5**

15388 UX codings in the DESCRIPTION, RETURN VALUE, and ERRORS sections have been changed
 15389 to EX.

15390 **Issue 6**

15391 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 15392 • The [ENAMETOOLONG] error is restored as an error dependent on _POSIX_NO_TRUNC.
 15393 This is since behavior may vary from one file system to another.

15394 The wording of the mandatory [ELOOP] error condition is updated, and a second optional
 15395 [ELOOP] error condition is added.

15396 **NAME**

15397 funlockfile — stdio locking functions

15398 **SYNOPSIS**

15399 TSF #include <stdio.h>

15400 void funlockfile(FILE *file);

15401

15402 **DESCRIPTION**

15403 Refer to *flockfile()*.

15404 **NAME**

15405 fwide — set stream orientation

15406 **SYNOPSIS**

15407 #include <stdio.h>

15408 #include <wchar.h>

15409 int fwide(FILE **stream*, int *mode*);15410 **DESCRIPTION**

15411 CX The functionality described on this reference page is aligned with the ISO C standard. Any
15412 conflict between the requirements described here and the ISO C standard is unintentional. This
15413 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

15414 The *fwide()* function shall determine the orientation of the stream pointed to by *stream*. If *mode* is
15415 greater than zero, the function first attempts to make the stream wide-oriented. If *mode* is less
15416 than zero, the function first attempts to make the stream byte-oriented. Otherwise, *mode* is zero
15417 and the function does not alter the orientation of the stream.

15418 If the orientation of the stream has already been determined, *fwide()* shall not change it.

15419 Because no return value is reserved to indicate an error, an application wishing to check for error
15420 situations should set *errno* to 0, then call *fwide()*, then check *errno*, and if it is non-zero, assume
15421 an error has occurred.

15422 **RETURN VALUE**

15423 The *fwide()* function shall return a value greater than zero if, after the call, the stream has wide-
15424 orientation, a value less than zero if the stream has byte-orientation, or zero if the stream has no
15425 orientation.

15426 **ERRORS**15427 The *fwide()* function may fail if:

15428 CX [EBADF] The *stream* argument is not a valid stream.

15429 **EXAMPLES**

15430 None.

15431 **APPLICATION USAGE**

15432 A call to *fwide()* with *mode* set to zero can be used to determine the current orientation of a
15433 stream.

15434 **RATIONALE**

15435 None.

15436 **FUTURE DIRECTIONS**

15437 None.

15438 **SEE ALSO**

15439 The Base Definitions volume of IEEE Std. 1003.1-200x, <wchar.h>

15440 **CHANGE HISTORY**

15441 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995
15442 (E).

15443 **Issue 6**

15444 Extensions beyond the ISO C standard are now marked.

15445 NAME

15446 fwprintf, swprintf, wprintf — print formatted wide-character output

15447 SYNOPSIS

15448 #include <stdio.h>

15449 #include <wchar.h>

15450 int fwprintf(FILE *restrict *stream*, const wchar_t *restrict *format*, ...);15451 int swprintf(wchar_t *restrict *ws*, size_t *n*, const wchar_t *restrict *format*, ...);15452 int wprintf(const wchar_t *restrict *format*, ...);

15453 DESCRIPTION

15454 CX The functionality described on this reference page is aligned with the ISO C standard. Any
 15455 conflict between the requirements described here and the ISO C standard is unintentional. This
 15456 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

15457 The *fwprintf()* function places output on the named output *stream*. The *wprintf()* function places
 15458 output on the standard output stream *stdout*. The *swprintf()* function places output followed by
 15459 the null wide character in consecutive wide characters starting at **ws*; no more than *n* wide
 15460 characters are written, including a terminating null wide character, which is always added
 15461 (unless *n* is zero).

15462 Each of these functions converts, formats, and prints its arguments under control of the *format*
 15463 wide-character string. The *format* is composed of zero or more directives: *ordinary wide-*
 15464 *characters*, which are simply copied to the output stream, and *conversion specifications*, each of
 15465 which results in the fetching of zero or more arguments. The results are undefined if there are
 15466 insufficient arguments for the *format*. If the *format* is exhausted while arguments remain, the
 15467 excess arguments are evaluated but are otherwise ignored.

15468 XSI Conversions can be applied to the *n*th argument after the *format* in the argument list, rather than
 15469 to the next unused argument. In this case, the conversion wide character '%' (see below) is
 15470 replaced by the sequence "%n\$", where *n* is a decimal integer in the range [1,{NL_ARGMAX}],
 15471 giving the position of the argument in the argument list. This feature provides for the definition
 15472 of *format* wide-character strings that select arguments in an order appropriate to specific
 15473 languages (see the EXAMPLES section).

15474 In *format* wide-character strings containing the "%n\$" form of conversion specifications,
 15475 numbered arguments in the argument list can be referenced from the *format* wide-character
 15476 string as many times as required.

15477 In *format* wide-character strings containing the '%' form of conversion specifications, each
 15478 argument in the argument list is used exactly once.

15479 All forms of the *fwprintf()* function allow for the insertion of a locale-dependent radix character
 15480 in the output string, output as a wide-character value. The radix character is defined in the
 15481 program's locale (category *LC_NUMERIC*). In the POSIX locale, or in a locale where the radix
 15482 character is not defined, the radix character defaults to a period ('.').

15483 XSI Each conversion specification is introduced by the '%' wide character or by the wide-character
 15484 sequence "%n\$", after which the following appear in sequence:

- 15485 • Zero or more *flags* (in any order), which modify the meaning of the conversion specification.
- 15486 • An optional minimum *field width*. If the converted value has fewer wide characters than the
 15487 field width, it shall be padded with spaces by default on the left; it shall be padded on the
 15488 right, if the left-adjustment flag ('-'), described below, is given to the field width. The field
 15489 width takes the form of an asterisk ('*'), described below, or a decimal integer.

- 15490 • An optional *precision* that gives the minimum number of digits to appear for the *d*, *i*, *o*, *u*, *x*,
 15491 and *X* conversions; the number of digits to appear after the radix character for the *e*, *E*, and *f*
 15492 conversions; the maximum number of significant digits for the *g* and *G* conversions; or the
 15493 maximum number of wide characters to be printed from a string in *s* conversions. The
 15494 precision takes the form of a period (*'.'*) followed either by an asterisk (*'*'*), described
 15495 below, or an optional decimal digit string, where a null digit string is treated as 0. If a
 15496 precision appears with any other conversion wide character, the behavior is undefined.
- 15497 • An optional length modifier that specifies the size of the argument.
- 15498 • A *conversion wide character* that indicates the type of conversion to be applied.
- 15499 A field width, or precision, or both, may be indicated by an asterisk (*'*'*). In this case an
 15500 argument of type `int` supplies the field width or precision. The application shall ensure that
 15501 arguments specifying field width, or precision, or both appear in that order before the argument,
 15502 if any, to be converted. A negative field width is taken as a *'-'* flag followed by a positive field
 15503 XSI width. A negative precision is taken as if the precision were omitted. In format wide-character
 15504 strings containing the "`%n$`" form of a conversion specification, a field width or precision may
 15505 be indicated by the sequence "`*m$`", where *m* is a decimal integer in the range
 15506 [1,{NL_ARGMAX}] giving the position in the argument list (after the format argument) of an
 15507 integer argument containing the field width or precision, for example:
- ```
15508 wprintf(L"%1$d:%2$.*3$d:%4$.*3$d\n", hour, min, precision, sec);
```
- 15509 The *format* can contain either numbered argument specifications (that is, "`%n$`" and "`*m$`"), or  
 15510 unnumbered argument specifications (that is, *'%'* and *'\*'*), but normally not both. The only  
 15511 exception to this is that "`%%`" can be mixed with the "`%n$`" form. The results of mixing  
 15512 numbered and unnumbered argument specifications in a *format* wide-character string are  
 15513 undefined. When numbered argument specifications are used, specifying the *N*th argument  
 15514 requires that all the leading arguments, from the first to the (*N*-1)th, are specified in the format  
 15515 wide-character string.
- 15516 The flag wide characters and their meanings are:
- 15517 XSI *'* The integer portion of the result of a decimal conversion (`%i`, `%d`, `%u`, `%f`, `%g`, or `%G`)  
 15518 shall be formatted with thousands' grouping wide characters. For other conversions,  
 15519 the behavior is undefined. The numeric grouping wide character is used.
- 15520 *-* The result of the conversion shall be left-justified within the field. The conversion shall  
 15521 be right-justified if this flag is not specified.
- 15522 *+* The result of a signed conversion shall always begin with a sign (*'+' or '-'*). The  
 15523 conversion shall begin with a sign only when a negative value is converted if this flag is  
 15524 not specified.
- 15525 *<space>* If the first wide character of a signed conversion is not a sign, or if a signed conversion  
 15526 results in no wide characters, a space shall be prefixed to the result. This means that if  
 15527 the *<space>* and *'+'* flags both appear, the space flag shall be ignored.
- 15528 *#* This flag specifies that the value is to be converted to an alternative form. For *o*  
 15529 conversion, it increases the precision (if necessary) to force the first digit of the result to  
 15530 be 0. For *x* or *X* conversions, a non-zero result shall have 0*x* (or 0*X*) prefixed to it. For *e*,  
 15531 *E*, *f*, *g*, or *G* conversions, the result shall always contain a radix character, even if no  
 15532 digits follow it. Without this flag, a radix character appears in the result of these  
 15533 conversions only if a digit follows it. For *g* and *G* conversions, trailing zeros shall *not* be  
 15534 removed from the result as they normally are. For other conversions, the behavior is  
 15535 undefined.

- 15536           0           For *d, i, o, u, x, X, e, E, f, g,* and *G* conversions, leading zeros (following any indication of sign or base) are used to pad to the field width; no space padding is performed. If the 0 and '-' flags both appear, the 0 flag shall be ignored. For *d, i, o, u, x,* and *X* conversions, if a precision is specified, the 0 flag shall be ignored. If the 0 and ''' flags both appear, the grouping wide characters are inserted before zero padding. For other conversions, the behavior is undefined.
- 15537
- 15538
- 15539
- 15540
- 15541
- 15542           The length modifiers and their meanings are:
- 15543           *hh*           Specifies that a following *d, i, o, u, x,* or *X* conversion specifier applies to a **signed char** or **unsigned char** argument (the argument will have been promoted according to the integer promotions, but its value shall be converted to **signed char** or **unsigned char** before printing); or that a following *n* conversion specifier applies to a pointer to a **signed char** argument.
- 15544
- 15545
- 15546
- 15547
- 15548           *h*            Specifies that a following *d, i, o, u, x,* or *X* conversion specifier applies to a **short** or **unsigned short** argument (the argument will have been promoted according to the integer promotions, but its value shall be converted to **short** or **unsigned short** before printing); or that a following *n* conversion specifier applies to a pointer to a **short** argument.
- 15549
- 15550
- 15551
- 15552
- 15553           *l* (ell)       Specifies that a following *d, i, o, u, x,* or *X* conversion specifier applies to a **long** or **unsigned long** argument; that a following *n* conversion specifier applies to a pointer to a **long** argument; that a following *c* conversion specifier applies to a **wint\_t** argument; that a following *s* conversion specifier applies to a pointer to a **wchar\_t** argument; or has no effect on a following *a, A, e, E, f, F, g,* or *G* conversion specifier.
- 15554
- 15555
- 15556
- 15557
- 15558           *ll* (ell-ell) Specifies that a following *d, i, o, u, x,* or *X* conversion specifier applies to a **long long** or **unsigned long long** argument; or that a following *n* conversion specifier applies to a pointer to a **long long** argument.
- 15559
- 15560
- 15561           *j*            Specifies that a following *d, i, o, u, x,* or *X* conversion specifier applies to an **intmax\_t** or **uintmax\_t** argument; or that a following *n* conversion specifier applies to a pointer to an **intmax\_t** argument.
- 15562
- 15563
- 15564           *z*            Specifies that a following *d, i, o, u, x,* or *X* conversion specifier applies to a **size\_t** or the corresponding signed integer type argument; or that a following *n* conversion specifier applies to a pointer to a signed integer type corresponding to **size\_t** argument.
- 15565
- 15566
- 15567           *t*            Specifies that a following *d, i, o, u, x,* or *X* conversion specifier applies to a **ptrdiff\_t** or the corresponding **unsigned** type argument; or that a following *n* conversion specifier applies to a pointer to a **ptrdiff\_t** argument.
- 15568
- 15569
- 15570           *L*            Specifies that a following *a, A, e, E, f, F, g,* or *G* conversion specifier applies to a **long double** argument.
- 15571
- 15572           If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.
- 15573
- 15574           The conversion wide characters and their meanings are:
- 15575           *d, i*          The **int** argument is converted to a signed decimal in the style `[-]dddd`. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it shall be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide characters.
- 15576
- 15577
- 15578
- 15579
- 15580           *o*            The **unsigned** argument is converted to unsigned octal format in the style `dddd`. The precision specifies the minimum number of digits to appear; if the value being
- 15581

|       |             |                                                                                                                               |
|-------|-------------|-------------------------------------------------------------------------------------------------------------------------------|
| 15582 |             | converted can be represented in fewer digits, it shall be expanded with leading zeros.                                        |
| 15583 |             | The default precision is 1. The result of converting 0 with an explicit precision of 0 is no                                  |
| 15584 |             | wide characters.                                                                                                              |
| 15585 | <i>u</i>    | The <b>unsigned</b> argument is converted to unsigned decimal format in the style <i>ddd</i> . The                            |
| 15586 |             | precision specifies the minimum number of digits to appear; if the value being                                                |
| 15587 |             | converted can be represented in fewer digits, it shall be expanded with leading zeros.                                        |
| 15588 |             | The default precision is 1. The result of converting 0 with an explicit precision of 0 is no                                  |
| 15589 |             | wide characters.                                                                                                              |
| 15590 | <i>x</i>    | The <b>unsigned</b> argument is converted to unsigned hexadecimal format in the style <i>ddd</i> ;                            |
| 15591 |             | the letters "abcdef" are used. The precision specifies the minimum number of digits                                           |
| 15592 |             | to appear; if the value being converted can be represented in fewer digits, it shall be                                       |
| 15593 |             | expanded with leading zeros. The default precision is 1. The result of converting 0 with                                      |
| 15594 |             | an explicit precision of 0 is no wide characters.                                                                             |
| 15595 | <i>X</i>    | Behaves the same as the <i>x</i> conversion wide character except that letters "ABCDEF" are                                   |
| 15596 |             | used instead of "abcdef".                                                                                                     |
| 15597 | <i>f, F</i> | The <b>double</b> argument is converted to decimal notation in the style <i>[-]ddd.ddd</i> , where                            |
| 15598 |             | the number of digits after the radix character is equal to the precision specification. If                                    |
| 15599 |             | the precision is missing, it is taken as 6; if the precision is explicitly 0 and no '#' flag is                               |
| 15600 |             | present, no radix character appears. If a radix character appears, at least one digit                                         |
| 15601 |             | appears before it. The value is rounded to the appropriate number of digits.                                                  |
| 15602 |             | A <b>double</b> argument representing an infinity is converted in one of the styles <i>[-]inf</i> or                          |
| 15603 |             | <i>[-]infinity</i> ; which style is implementation-defined. A <b>double</b> argument representing a                           |
| 15604 |             | NaN is converted in one of the styles <i>[-]nan</i> or <i>[-]nan(n-char-sequence)</i> ; which style, and                      |
| 15605 |             | the meaning of any <i>n-char-sequence</i> , is implementation-defined. The <i>F</i> conversion                                |
| 15606 |             | specifier produces INF, INFINITY, or NAN instead of inf, infinity, or nan, respectively.                                      |
| 15607 | <i>e, E</i> | The <b>double</b> argument is converted in the style <i>[-]d.ddde±dd</i> , where there is one digit                           |
| 15608 |             | before the radix character (which is non-zero if the argument is non-zero) and the                                            |
| 15609 |             | number of digits after it is equal to the precision; if the precision is missing, it is taken                                 |
| 15610 |             | as 6; if the precision is 0 and no '#' flag is present, no radix character appears. The                                       |
| 15611 |             | value is rounded to the appropriate number of digits. The <i>E</i> conversion wide character                                  |
| 15612 |             | shall produce a number with <i>E</i> instead of <i>e</i> introducing the exponent. The exponent                               |
| 15613 |             | always contains at least two digits. If the value is 0, the exponent is 0.                                                    |
| 15614 |             | A <b>double</b> argument representing an infinity or NaN is converted in the style of an <i>f</i> or <i>F</i>                 |
| 15615 |             | conversion specifier.                                                                                                         |
| 15616 | <i>g, G</i> | The <b>double</b> argument is converted in the style <i>f</i> or <i>e</i> (or in the style <i>E</i> in the case of a <i>G</i> |
| 15617 |             | conversion wide character), with the precision specifying the number of significant                                           |
| 15618 |             | digits. If an explicit precision is 0, it is taken as 1. The style used depends on the value                                  |
| 15619 |             | converted; style <i>e</i> (or <i>E</i> ) shall be used only if the exponent resulting from such a                             |
| 15620 |             | conversion is less than -4 or greater than or equal to the precision. Trailing zeros are                                      |
| 15621 |             | removed from the fractional portion of the result; a radix character appears only if it is                                    |
| 15622 |             | followed by a digit.                                                                                                          |
| 15623 |             | A <b>double</b> argument representing an infinity or NaN is converted in the style of an <i>f</i> or <i>F</i>                 |
| 15624 |             | conversion specifier.                                                                                                         |
| 15625 | <i>a, A</i> | A <b>double</b> argument representing a floating-point number is converted in the style                                       |
| 15626 |             | <i>[-]0xh.hhhh p1d</i> , where there is one hexadecimal digit (which is non-zero if the                                       |
| 15627 |             | argument is a normalized floating-point number and is otherwise unspecified) before                                           |
| 15628 |             | the decimal-point character 235) and the number of hexadecimal digits after it is equal                                       |

15629 to the precision; if the precision is missing and FLT\_RADIX is a power of 2, then the  
 15630 precision is sufficient for an exact representation of the value; if the precision is missing  
 15631 and FLT\_RADIX is not a power of 2, then the precision is sufficient to distinguish 236  
 15632 values of type **double**, except that trailing zeros may be omitted; if the precision is zero  
 15633 and the '#' flag is not specified, no decimal-point character appears. The letters  
 15634 "abcdef" are used for a conversion and the letters "ABCDEF" for A conversion. The A  
 15635 conversion specifier produces a number with 'X' and 'P' instead of 'x' and 'p'.  
 15636 The exponent always contains at least one digit, and only as many more digits as  
 15637 necessary to represent the decimal exponent of 2. If the value is zero, the exponent is  
 15638 zero.

15639 A **double** argument representing an infinity or NaN is converted in the style of an *f* or *F*  
 15640 conversion specifier.

15641 *c* If no *l* (ell) qualifier is present, the **int** argument is converted to a wide character as if by  
 15642 calling the *btowc*() function and the resulting wide character is written. Otherwise, the  
 15643 **wint\_t** argument is converted to **wchar\_t**, and written.

15644 *s* If no *l* (ell) qualifier is present, the application shall ensure that the argument is a  
 15645 pointer to a character array containing a character sequence beginning in the initial  
 15646 shift state. Characters from the array are converted as if by repeated calls to the  
 15647 *mbrtowc*() function, with the conversion state described by an **mbstate\_t** object  
 15648 initialized to zero before the first character is converted, and written up to (but not  
 15649 including) the terminating null wide character. If the precision is specified, no more  
 15650 than that many wide characters are written. If the precision is not specified, or is  
 15651 greater than the size of the array, the application shall ensure that the array contains a  
 15652 null wide character.

15653 If an *l* (ell) qualifier is present, the application shall ensure that the argument is a  
 15654 pointer to an array of type **wchar\_t**. Wide characters from the array are written up to  
 15655 (but not including) a terminating null wide character. If no precision is specified, or is  
 15656 greater than the size of the array, the application shall ensure that the array contains a  
 15657 null wide character. If a precision is specified, no more than that many wide characters  
 15658 are written.

15659 *p* The application shall ensure that the argument is a pointer to **void**. The value of the  
 15660 pointer is converted to a sequence of printable wide characters, in an implementation-  
 15661 defined manner.

15662 *n* The application shall ensure that the argument is a pointer to an integer into which is  
 15663 written the number of wide characters written to the output so far by this call to one of  
 15664 the *fwprintf*() functions. No argument is converted, but one is consumed. If the  
 15665 conversion specification includes any flags, a field width, or a precision, the behavior is  
 15666 undefined.

15667 XSI *C* Same as *lc*.

15668 XSI *S* Same as *ls*.

15669 % Output a '%' wide character; no argument is converted. The entire conversion  
 15670 specification shall be "%%".

15671 If a conversion specification does not match one of the above forms, the behavior is undefined.

15672 In no case does a nonexistent or small field width cause truncation of a field; if the result of a  
 15673 conversion is wider than the field width, the field is simply expanded to contain the conversion  
 15674 result. Characters generated by *fwprintf*() and *wprintf*() are printed as if *fputwc*() had been  
 15675 called.

15676 For *a* and *A* conversions, if `FLT_RADIX` is a power of 2, the value is correctly rounded to a  
15677 hexadecimal floating number with the given precision.

15678 If `FLT_RADIX` is not a power of 2, the result should be one of the two adjacent numbers in  
15679 hexadecimal floating style with the given precision, with the extra stipulation that the error  
15680 should have a correct sign for the current rounding direction.

15681 For *e*, *E*, *f*, *F*, *g*, and *G* conversions, if the number of significant decimal digits is at most  
15682 `DECIMAL_DIG`, then the result should be correctly rounded. If the number of significant  
15683 decimal digits is more than `DECIMAL_DIG` but the source value is exactly representable with  
15684 `DECIMAL_DIG` digits, then the result should be an exact representation with trailing zeros.  
15685 Otherwise, the source value is bounded by two adjacent decimal strings "*L* < *U*", both having  
15686 `DECIMAL_DIG` significant digits; the value of the resultant decimal string "*D*" should satisfy "*L*  
15687 <= *D* <= *U*", with the extra stipulation that the error should have a correct sign for the current  
15688 rounding direction.

15689 CX The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the call to a  
15690 successful execution of `fwprintf()` or `wprintf()` and the next successful completion of a call to  
15691 `fflush()` or `fclose()` on the same stream, or a call to `exit()` or `abort()`.

#### 15692 RETURN VALUE

15693 Upon successful completion, these functions shall return the number of wide characters  
15694 transmitted, excluding the terminating null wide character in the case of `swprintf()`, or a negative  
15695 CX value if an output error was encountered, and set *errno* to indicate the error.

15696 If *n* or more wide characters were requested to be written, `swprintf()` shall return a negative  
15697 CX value, and set *errno* to indicate the error.

#### 15698 ERRORS

15699 For the conditions under which `fwprintf()` and `wprintf()` fail and may fail, refer to `fputwc()`.

15700 In addition, all forms of `wprintf()` may fail if:

15701 XSI `[EILSEQ]` A wide-character code that does not correspond to a valid character has been  
15702 detected.

15703 XSI `[EINVAL]` There are insufficient arguments.

15704 In addition, `wprintf()` and `fwprintf()` may fail if:

15705 XSI `[ENOMEM]` Insufficient storage space is available.

#### 15706 EXAMPLES

15707 To print the language-independent date and time format, the following statement could be used:

```
15708 wprintf(format, weekday, month, day, hour, min);
```

15709 For American usage, *format* could be a pointer to the wide-character string:

```
15710 L"%s, %s %d, %d:%.2d\n"
```

15711 producing the message:

```
15712 Sunday, July 3, 10:02
```

15713 whereas for German usage, *format* could be a pointer to the wide-character string:

```
15714 L"%1$s, %3$d. %2$s, %4$d:%5$.2d\n"
```

15715 producing the message:

```
15716 Sonntag, 3. Juli, 10:02
```

15717 **APPLICATION USAGE**

15718 None.

15719 **RATIONALE**

15720 None.

15721 **FUTURE DIRECTIONS**

15722 None.

15723 **SEE ALSO**

15724 *btowc()*, *fputwc()*, *fwscanf()*, *mbrtowc()*, *setlocale()*, the Base Definitions volume of  
15725 IEEE Std. 1003.1-200x, `<stdio.h>`, `<wchar.h>`, the Base Definitions volume of  
15726 IEEE Std. 1003.1-200x, Chapter 7, Locale

15727 **CHANGE HISTORY**

15728 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
15729 (E).

15730 **Issue 6**

15731 The Open Group corrigenda item U040/1 has been applied to the RETURN VALUE section,  
15732 describing the case if *n* or more wide characters are requested to be written using *swprintf()*.

15733 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

15734 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 15735 • The prototypes for *fwprintf()*, *swprintf()*, and *wprintf()* are updated.
- 15736 • The DESCRIPTION is updated.

15737 **NAME**

15738 fwrite — binary output

15739 **SYNOPSIS**

15740 #include &lt;stdio.h&gt;

15741 size\_t fwrite(const void \*restrict ptr, size\_t size, size\_t nitems,  
15742 FILE \*restrict stream);15743 **DESCRIPTION**15744 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
15745 conflict between the requirements described here and the ISO C standard is unintentional. This  
15746 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.15747 The *fwrite()* function shall write, from the array pointed to by *ptr*, up to *nitems* members whose  
15748 size is specified by *size*, to the stream pointed to by *stream*. For each object, *size* calls are made to  
15749 the *fputc()* function, taking the values (in order) from an array of **unsigned char** exactly  
15750 overlaying the object. The file-position indicator for the stream (if defined) shall be advanced by  
15751 the number of bytes successfully written. If an error occurs, the resulting value of the file-  
15752 position indicator for the stream is indeterminate.15753 CX The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the successful  
15754 execution of *fwrite()* and the next successful completion of a call to *fflush()* or *fclose()* on the  
15755 same stream, or a call to *exit()* or *abort()*.15756 **RETURN VALUE**15757 The *fwrite()* function shall return the number of members successfully written, which may be  
15758 less than *nitems* if a write error is encountered. If *size* or *nitems* is 0, *fwrite()* shall return 0 and the  
15759 state of the stream remains unchanged. Otherwise, if a write error occurs, the error indicator for  
15760 CX the stream shall be set, and *errno* shall be set to indicate the error.15761 **ERRORS**15762 Refer to *fputc()*.15763 **EXAMPLES**

15764 None.

15765 **APPLICATION USAGE**15766 Because of possible differences in member length and byte ordering, files written using *fwrite()*  
15767 are application-dependent, and possibly cannot be read using *fread()* by a different application  
15768 or by the same application on a different processor.15769 **RATIONALE**

15770 None.

15771 **FUTURE DIRECTIONS**

15772 None.

15773 **SEE ALSO**15774 *ferror()*, *fopen()*, *printf()*, *putc()*, *puts()*, *write()*, the Base Definitions volume of  
15775 IEEE Std. 1003.1-200x, <stdio.h>15776 **CHANGE HISTORY**

15777 First released in Issue 1. Derived from Issue 1 of the SVID.

15778 **Issue 4**15779 In the DESCRIPTION, the text is changed to make it clear that the function advances the file-  
15780 position indicator by the number of bytes successfully written rather than the number of  
15781 characters, which could include multi-byte sequences.

15782 The following change is incorporated for alignment with the ISO C standard:

- 15783
- The type of argument *ptr* is changed from **void\*** to **const void\***.

15784 **Issue 6**

15785 Extensions beyond the ISO C standard are now marked.

15786 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 15787
- The *fwrite()* prototype is updated.
- 15788
- The DESCRIPTION is updated to clarify how the data is written out using *fputc()*.

15789 **NAME**

15790 fwscanf, swscanf, wscanf — convert formatted wide-character input

15791 **SYNOPSIS**

15792 #include &lt;stdio.h&gt;

15793 #include &lt;wchar.h&gt;

15794 int fwscanf(FILE \*restrict *stream*, const wchar\_t \*restrict *format*, ... );15795 int swscanf(const wchar\_t \*restrict *ws*,15796 const wchar\_t \*restrict *format*, ... );15797 int wscanf(const wchar\_t \**format*, ... );15798 **DESCRIPTION**

15799 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 15800 conflict between the requirements described here and the ISO C standard is unintentional. This  
 15801 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

15802 The *fwscanf()* function reads from the named input *stream*. The *wscanf()* function reads from the  
 15803 standard input stream *stdin*. The *swscanf()* function reads from the wide-character string *ws*.  
 15804 Each function reads wide characters, interprets them according to a format, and stores the  
 15805 results in its arguments. Each expects, as arguments, a control wide-character string *format*  
 15806 described below, and a set of *pointer* arguments indicating where the converted input should be  
 15807 stored. The result is undefined if there are insufficient arguments for the format. If the format is  
 15808 exhausted while arguments remain, the excess arguments are evaluated but are otherwise  
 15809 ignored.

15810 XSI Conversions can be applied to the *n*th argument after the *format* in the argument list, rather than  
 15811 to the next unused argument. In this case, the conversion wide character '%' (see below) is  
 15812 replaced by the sequence "%n\$", where *n* is a decimal integer in the range [1,{NL\_ARGMAX}].  
 15813 This feature provides for the definition of format wide-character strings that select arguments in  
 15814 an order appropriate to specific languages. In format wide-character strings containing the  
 15815 "%n\$" form of conversion specifications, it is unspecified whether numbered arguments in the  
 15816 argument list can be referenced from the format wide-character string more than once.

15817 The *format* can contain either form of a conversion specification—that is, '%' or "%n\$"—but the  
 15818 two forms cannot normally be mixed within a single *format* wide-character string. The only  
 15819 exception to this is that "%%" or "%\*" can be mixed with the "%n\$" form.

15820 The *fwscanf()* function in all its forms allows for detection of a language-dependent radix  
 15821 character in the input string, encoded as a wide-character value. The radix character is defined in  
 15822 the program's locale (category *LC\_NUMERIC*). In the POSIX locale, or in a locale where the  
 15823 radix character is not defined, the radix character defaults to a period ('.').

15824 The *format* is a wide-character string composed of zero or more directives. Each directive is  
 15825 composed of one of the following: one or more white-space wide characters (<space>, <tab>,  
 15826 <newline>, <vertical-tab>, or <form-feed> characters); an ordinary wide character (neither '%'   
 15827 nor a white-space character); or a conversion specification. Each conversion specification is  
 15828 XSI introduced by a '%' or the sequence "%n\$" after which the following appear in sequence:

- 15829 • An optional assignment-suppressing character '\*'.
- 15830 • An optional non-zero decimal integer that specifies the maximum field width.
- 15831 • An optional length modifier that specifies the size of the receiving object.
- 15832 • A conversion wide character that specifies the type of conversion to be applied. The valid  
 15833 conversion wide characters are described below.

15834 The *fwscanf()* functions execute each directive of the format in turn. If a directive fails, as  
 15835 detailed below, the function shall return. Failures are described as input failures (due to the  
 15836 unavailability of input bytes) or matching failures (due to inappropriate input).

15837 A directive composed of one or more white-space wide characters is executed by reading input  
 15838 until no more valid input can be read, or up to the first wide character which is not a white-  
 15839 space wide character, which remains unread.

15840 A directive that is an ordinary wide character is executed as follows. The next wide character is  
 15841 read from the input and compared with the wide character that comprises the directive; if the  
 15842 comparison shows that they are not equivalent, the directive fails, and the differing and  
 15843 subsequent wide characters remain unread. Similarly, if end-of-file, an encoding error, or a read  
 15844 error prevents a wide character from being read, the directive fails.

15845 A directive that is a conversion specification defines a set of matching input sequences, as  
 15846 described below for each conversion wide character. A conversion specification is executed in  
 15847 the following steps.

15848 Input white-space wide characters (as specified by *iswspace()*) are skipped, unless the conversion  
 15849 specification includes a ' [ ', c, or n conversion character.

15850 An item is read from the input, unless the conversion specification includes an n conversion  
 15851 wide character. An input item is defined as the longest sequence of input wide characters, not  
 15852 exceeding any specified field width, which is an initial subsequence of a matching sequence. The  
 15853 first wide character, if any, after the input item remains unread. If the length of the input item is  
 15854 0, the execution of the conversion specification fails; this condition is a matching failure, unless  
 15855 end-of-file, an encoding error, or a read error prevented input from the stream, in which case it is  
 15856 an input failure.

15857 Except in the case of a ' % ' conversion wide character, the input item (or, in the case of a %n  
 15858 conversion specification, the count of input wide characters) is converted to a type appropriate  
 15859 to the conversion wide character. If the input item is not a matching sequence, the execution of  
 15860 the conversion specification fails; this condition is a matching failure. Unless assignment  
 15861 suppression was indicated by a ' \* ', the result of the conversion is placed in the object pointed  
 15862 to by the first argument following the *format* argument that has not already received a  
 15863 XSI conversion result if the conversion specification is introduced by ' % ', or in the *n*th argument if  
 15864 introduced by the wide-character sequence "%n\$". If this object does not have an appropriate  
 15865 type, or if the result of the conversion cannot be represented in the space provided, the behavior  
 15866 is undefined.

15867 The length modifiers and their meanings are:

15868 *hh* Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument  
 15869 with type pointer to **signed char** or **unsigned char**.

15870 *h* Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument  
 15871 with type pointer to **short** or **unsigned short**.

15872 *l* (ell) Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument  
 15873 with type pointer to **long** or **unsigned long**; that a following *a, A, e, E, f, F, g,* or *G*  
 15874 conversion specifier applies to an argument with type pointer to **double**; or that a  
 15875 following *c, s,* or ' [ ' conversion specifier applies to an argument with type pointer to  
 15876 **wchar\_t**.

15877 *ll* (ell-ell) Specifies that a following *d, i, o, u, x, X,* or *n* conversion specifier applies to an argument  
 15878 with type pointer to **long long** or **unsigned long long**.

|       |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 15879 | <i>j</i>          | Specifies that a following <i>d, i, o, u, x, X,</i> or <i>n</i> conversion specifier applies to an argument with type pointer to <b>intmax_t</b> or <b>uintmax_t</b> .                                                                                                                                                                                                                                                                                                                                                                                   |
| 15880 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15881 | <i>z</i>          | Specifies that a following <i>d, i, o, u, x, X,</i> or <i>n</i> conversion specifier applies to an argument with type pointer to <b>size_t</b> or the corresponding signed integer type.                                                                                                                                                                                                                                                                                                                                                                 |
| 15882 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15883 | <i>t</i>          | Specifies that a following <i>d, i, o, u, x, X,</i> or <i>n</i> conversion specifier applies to an argument with type pointer to <b>ptrdiff_t</b> or the corresponding <b>unsigned</b> type.                                                                                                                                                                                                                                                                                                                                                             |
| 15884 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15885 | <i>L</i>          | Specifies that a following <i>a, A, e, E, f, F, g,</i> or <i>G</i> conversion specifier applies to an argument with type pointer to <b>long double</b> .                                                                                                                                                                                                                                                                                                                                                                                                 |
| 15886 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15887 |                   | If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 15888 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15889 |                   | The following conversion wide characters are valid:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 15890 | <i>d</i>          | Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of <i>wcstol()</i> with the value 10 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to <b>int</b> .                                                                                                                                                                                                                                              |
| 15891 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15892 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15893 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15894 | <i>i</i>          | Matches an optionally signed integer, whose format is the same as expected for the subject sequence of <i>wcstol()</i> with 0 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to <b>int</b> .                                                                                                                                                                                                                                                                 |
| 15895 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15896 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15897 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15898 | <i>o</i>          | Matches an optionally signed octal integer, whose format is the same as expected for the subject sequence of <i>wcstoul()</i> with the value 8 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to <b>unsigned</b> .                                                                                                                                                                                                                                           |
| 15899 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15900 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15901 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15902 | <i>u</i>          | Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of <i>wcstoul()</i> with the value 10 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to <b>unsigned</b> .                                                                                                                                                                                                                                        |
| 15903 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15904 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15905 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15906 | <i>x</i>          | Matches an optionally signed hexadecimal integer, whose format is the same as expected for the subject sequence of <i>wcstoul()</i> with the value 16 for the <i>base</i> argument. In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to <b>unsigned</b> .                                                                                                                                                                                                                                    |
| 15907 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15908 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15909 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15910 | <i>a, e, f, g</i> | Matches an optionally signed floating-point number, infinity, or NaN whose format is the same as expected for the subject sequence of <i>wcstod()</i> . In the absence of a size modifier, the application shall ensure that the corresponding argument is a pointer to <b>float</b> .                                                                                                                                                                                                                                                                   |
| 15911 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15912 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15913 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15914 |                   | If the <i>fwprintf()</i> family of functions generates character string representations for infinity and NaN (a symbolic entity encoded in floating-point format) to support IEEE Std. 754-1985, the <i>fwscanf()</i> family of functions shall recognize them as input.                                                                                                                                                                                                                                                                                 |
| 15915 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15916 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15917 | <i>s</i>          | Matches a sequence of non white-space wide characters. If no <i>l</i> (ell) qualifier is present, characters from the input field are converted as if by repeated calls to the <i>wcrtomb()</i> function, with the conversion state described by an <b>mbstate_t</b> object initialized to zero before the first wide character is converted. The application shall ensure that the corresponding argument is a pointer to a character array large enough to accept the sequence and the terminating null character, which shall be added automatically. |
| 15918 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15919 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15920 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15921 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 15922 |                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

|       |   |                                                                                                             |
|-------|---|-------------------------------------------------------------------------------------------------------------|
| 15923 |   | Otherwise, the application shall ensure that the corresponding argument is a pointer to                     |
| 15924 |   | an array of <b>wchar_t</b> large enough to accept the sequence and the terminating null wide                |
| 15925 |   | character, which shall be added automatically.                                                              |
| 15926 | [ | Matches a non-empty sequence of wide characters from a set of expected wide                                 |
| 15927 |   | characters (the <i>scanset</i> ). If no <i>l</i> (ell) qualifier is present, wide characters from the input |
| 15928 |   | field are converted as if by repeated calls to the <i>wcrtomb()</i> function, with the conversion           |
| 15929 |   | state described by an <b>mbstate_t</b> object initialized to zero before the first wide character           |
| 15930 |   | is converted. The application shall ensure that the corresponding argument is a pointer                     |
| 15931 |   | to a character array large enough to accept the sequence and the terminating null                           |
| 15932 |   | character, which shall be added automatically.                                                              |
| 15933 |   | If an <i>l</i> (ell) qualifier is present, the application shall ensure that the corresponding              |
| 15934 |   | argument is a pointer to an array of <b>wchar_t</b> large enough to accept the sequence and                 |
| 15935 |   | the terminating null wide character, which shall be added automatically.                                    |
| 15936 |   | The conversion specification includes all subsequent wide characters in the <i>format</i>                   |
| 15937 |   | string up to and including the matching right square bracket (']'). The wide                                |
| 15938 |   | characters between the square brackets (the <i>scanlist</i> ) comprise the scanset, unless the              |
| 15939 |   | wide character after the left square bracket is a circumflex ('^'), in which case the                       |
| 15940 |   | scanset contains all wide characters that do not appear in the scanlist between the                         |
| 15941 |   | circumflex and the right square bracket. If the conversion specification begins with                        |
| 15942 |   | "[ ]" or "[ ^ ]", the right square bracket is included in the scanlist and the next right                   |
| 15943 |   | square bracket is the matching right square bracket that ends the conversion                                |
| 15944 |   | specification; otherwise, the first right square bracket is the one that ends the                           |
| 15945 |   | conversion specification. If a '-' is in the scanlist and is not the first wide character,                  |
| 15946 |   | nor the second where the first wide character is a '^', nor the last wide character, the                    |
| 15947 |   | behavior is implementation-defined.                                                                         |
| 15948 | c | Matches a sequence of wide characters of exactly the number specified by the field                          |
| 15949 |   | width (1 if no field width is present in the conversion specification).                                     |
| 15950 |   | If no <i>l</i> (ell) length modifier is present, characters from the input field are converted as if        |
| 15951 |   | by repeated calls to the <i>wcrtomb()</i> function, with the conversion state described by an               |
| 15952 |   | <b>mbstate_t</b> object initialized to zero before the first wide character is converted. The               |
| 15953 |   | corresponding argument shall be a pointer to the initial element of a character array                       |
| 15954 |   | large enough to accept the sequence. No null character is added.                                            |
| 15955 |   | If an <i>l</i> (ell) length modifier is present, the corresponding argument shall be a pointer to           |
| 15956 |   | the initial element of an array of <b>wchar_t</b> large enough to accept the sequence. No null              |
| 15957 |   | wide character is added.                                                                                    |
| 15958 |   | Otherwise, the application shall ensure that the corresponding argument is a pointer to                     |
| 15959 |   | an array of <b>wchar_t</b> large enough to accept the sequence. No null wide character is                   |
| 15960 |   | added.                                                                                                      |
| 15961 | p | Matches an implementation-defined set of sequences, which shall be the same as the set                      |
| 15962 |   | of sequences that is produced by the <i>%p</i> conversion of the corresponding <i>fwprintf()</i>            |
| 15963 |   | functions. The application shall ensure that the corresponding argument is a pointer to                     |
| 15964 |   | a pointer to <b>void</b> . The interpretation of the input item is implementation-defined. If the           |
| 15965 |   | input item is a value converted earlier during the same program execution, the pointer                      |
| 15966 |   | that results shall compare equal to that value; otherwise, the behavior of the <i>%p</i>                    |
| 15967 |   | conversion is undefined.                                                                                    |
| 15968 | n | No input is consumed. The application shall ensure that the corresponding argument is                       |
| 15969 |   | a pointer to the integer into which is to be written the number of wide characters read                     |
| 15970 |   | from the input so far by this call to the <i>fwscanf()</i> functions. Execution of a <i>%n</i>              |

- 15971 conversion specification does not increment the assignment count returned at the  
 15972 completion of execution of the function. No argument is converted, but one is  
 15973 consumed. If the conversion specification includes an assignment-suppressing wide  
 15974 character or a field width, the behavior is undefined.
- 15975 XSI **C** Same as *lc*.
- 15976 **S** Same as *ls*.
- 15977 % Matches a single '%'; no conversion or assignment occurs. The complete conversion  
 15978 specification shall be "%%".
- 15979 If a conversion specification is invalid, the behavior is undefined.
- 15980 The conversion characters *A*, *E*, *F*, *G*, and *X* are also valid and behave the same as, respectively, *a*,  
 15981 *e*, *f*, *g*, and *x*.
- 15982 If end-of-file is encountered during input, conversion is terminated. If end-of-file occurs before  
 15983 any wide characters matching the current conversion specification (except for %*n*) have been  
 15984 read (other than leading white-space, where permitted), execution of the current conversion  
 15985 specification terminates with an input failure. Otherwise, unless execution of the current  
 15986 conversion specification is terminated with a matching failure, execution of the following  
 15987 conversion specification (if any) is terminated with an input failure.
- 15988 Reaching the end of the string in *swscanf()* is equivalent to encountering end-of-file for *fwscanf()*.
- 15989 If conversion terminates on a conflicting input, the offending input is left unread in the input.  
 15990 Any trailing white space (including <newline>) is left unread unless matched by a conversion  
 15991 specification. The success of literal matches and suppressed assignments is only directly  
 15992 determinable via the %*n* conversion specification.
- 15993 The *fwscanf()* and *wscanf()* functions may mark the *st\_atime* field of the file associated with  
 15994 *stream* for update. The *st\_atime* field shall be marked for update by the first successful execution  
 15995 of *fgetc()*, *fgetwc()*, *fgets()*, *fgetws()*, *fread()*, *getc()*, *getwc()*, *getchar()*, *getwchar()*, *gets()*, *fscanf()*,  
 15996 or *fwscanf()* using *stream* that returns data not supplied by a prior call to *ungetc()*.
- 15997 **RETURN VALUE**
- 15998 Upon successful completion, these functions shall return the number of successfully matched  
 15999 and assigned input items; this number can be 0 in the event of an early matching failure. If the  
 16000 input ends before the first matching failure or conversion, EOF shall be returned. If a read error  
 16001 CX occurs the error indicator for the stream is set, EOF shall be returned, and *errno* shall be set to  
 16002 indicate the error.
- 16003 **ERRORS**
- 16004 For the conditions under which the *fwscanf()* functions shall fail and may fail, refer to *fgetwc()*.
- 16005 In addition, *fwscanf()* may fail if:
- 16006 XSI **[EILSEQ]** Input byte sequence does not form a valid character.
- 16007 XSI **[EINVAL]** There are insufficient arguments.

16008 **EXAMPLES**

16009 The call:

```
16010 int i, n; float x; char name[50];
16011 n = wscanf(L"%d%f%s", &i, &x, name);
```

16012 with the input line:

16013 25 54.32E-1 Hamster

16014 assigns to *n* the value 3, to *i* the value 25, to *x* the value 5.432, and *name* contains the string  
16015 "Hamster".

16016 The call:

```
16017 int i; float x; char name[50];
16018 (void) wscanf(L"%2d%f*d %[0123456789]", &i, &x, name);
```

16019 with input:

16020 56789 0123 56a72

16021 assigns 56 to *i*, 789.0 to *x*, skip 0123, and place the string "56\0" in *name*. The next call to  
16022 *getchar()* shall return the character 'a'.

16023 **APPLICATION USAGE**

16024 In format strings containing the '%' form of conversion specifications, each argument in the  
16025 argument list is used exactly once.

16026 **RATIONALE**

16027 None.

16028 **FUTURE DIRECTIONS**

16029 None.

16030 **SEE ALSO**

16031 *getwc()*, *fwprintf()*, *setlocale()*, *wcstod()*, *wcstol()*, *wcstoul()*, *wcrtomb()*, the Base Definitions  
16032 volume of IEEE Std. 1003.1-200x, <**langinfo.h**>, <**stdio.h**>, <**wchar.h**>, the Base Definitions  
16033 volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

16034 **CHANGE HISTORY**

16035 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
16036 (E).

16037 **Issue 6**

16038 The DESCRIPTION is updated to avoid use of the term "must" for application requirements.

16039 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 16040 • The prototypes for *fwscanf()* and *swscanf()* are updated.
- 16041 • The DESCRIPTION is updated.

16042 **NAME**

16043        *gai\_strerror* — address and name information error description

16044 **SYNOPSIS**

16045        #include <netdb.h>

16046        char \**gai\_strerror*(int *ecode*);

16047 **DESCRIPTION**

16048        The *gai\_strerror*() function shall return a text string describing an error that is listed in the  
16049        <netdb.h> header.

16050        When the *ecode* argument is one of the values listed in the <netdb.h> header, the function return  
16051        value points to a string describing the error. If the argument is not one of those values, the  
16052        function shall return a pointer to a string whose contents indicate an unknown error.

16053 **RETURN VALUE**

16054        Upon successful completion, *gai\_strerror*() shall return a pointer to an implementation-defined  
16055        string.

16056 **ERRORS**

16057        No errors are defined.

16058 **EXAMPLES**

16059        None.

16060 **APPLICATION USAGE**

16061        None.

16062 **RATIONALE**

16063        None.

16064 **FUTURE DIRECTIONS**

16065        None.

16066 **SEE ALSO**

16067        *getaddrinfo*(), the Base Definitions volume of IEEE Std. 1003.1-200x, <netdb.h>

16068 **CHANGE HISTORY**

16069        First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

16070 **NAME**

16071 gcvt — convert a floating-point number to a string (**LEGACY**)

16072 **SYNOPSIS**

16073 XSI #include <stdlib.h>

16074 char \*gcvt(double value, int ndigit, char \*buf);

16075

16076 **DESCRIPTION**

16077 Refer to *ecvt()*.

16078 **NAME**

16079           getaddrinfo — get address information

16080 **SYNOPSIS**

16081           #include &lt;sys/socket.h&gt;

16082           #include &lt;netdb.h&gt;

16083           int getaddrinfo(const char \*restrict *nodename*,16084                           const char \*restrict *servname*,16085                           const struct addrinfo \*restrict *hints*,16086                           struct addrinfo \*\*restrict *res*);16087 **DESCRIPTION**16088           Refer to *freeaddrinfo()*.

16089 **NAME**

16090        getc — get a byte from a stream

16091 **SYNOPSIS**

16092        #include &lt;stdio.h&gt;

16093        int getc(FILE \*stream);

16094 **DESCRIPTION**

16095 cx        The functionality described on this reference page is aligned with the ISO C standard. Any  
16096        conflict between the requirements described here and the ISO C standard is unintentional. This  
16097        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

16098        The *getc()* function shall be equivalent to *fgetc()*, except that if it is implemented as a macro it  
16099        may evaluate *stream* more than once, so the argument should never be an expression with side  
16100        effects.

16101 **RETURN VALUE**16102        Refer to *fgetc()*.16103 **ERRORS**16104        Refer to *fgetc()*.16105 **EXAMPLES**

16106        None.

16107 **APPLICATION USAGE**

16108        If the integer value returned by *getc()* is stored into a variable of type **char** and then compared  
16109        against the integer constant EOF, the comparison may never succeed, because sign-extension of  
16110        a variable of type **char** on widening to integer is implementation-defined.

16111        Because it may be implemented as a macro, *getc()* may treat incorrectly a *stream* argument with  
16112        side effects. In particular, *getc(\*f++)* does not necessarily work as expected. Therefore, use of this  
16113        function should be preceded by "#undef getc" in such situations; *fgetc()* could also be used.

16114 **RATIONALE**

16115        None.

16116 **FUTURE DIRECTIONS**

16117        None.

16118 **SEE ALSO**16119        *fgetc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>16120 **CHANGE HISTORY**

16121        First released in Issue 1. Derived from Issue 1 of the SVID.

16122 **Issue 4**

16123        The words “a character variable” are replaced by “a variable of type **char**”, to emphasize the fact  
16124        that this function deals with byte values.

16125        The APPLICATION USAGE section now states that the use of this function is not recommended.

16126 **NAME**

16127        getc\_unlocked, getchar\_unlocked, putc\_unlocked, putchar\_unlocked — stdio with explicit client  
16128        locking

16129 **SYNOPSIS**

```
16130 TSF #include <stdio.h>

16131 int getc_unlocked(FILE *stream);
16132 int getchar_unlocked(void);
16133 int putc_unlocked(int c, FILE *stream);
16134 int putchar_unlocked(int c);
16135
```

16136 **DESCRIPTION**

16137        Versions of the functions *getc()*, *getchar()*, *putc()*, and *putchar()* respectively named  
16138        *getc\_unlocked()*, *getchar\_unlocked()*, *putc\_unlocked()*, and *putchar\_unlocked()* shall be provided  
16139        which are functionally identical to the original versions, with the exception that they are not  
16140        required to be implemented in a thread-safe manner. They may only safely be used within a  
16141        scope protected by *flockfile()* (or *ftrylockfile()*) and *funlockfile()*. These functions may safely be  
16142        used in a multi-threaded program if and only if they are called while the invoking thread owns  
16143        the (**FILE\***) object, as is the case after a successful call of the *flockfile()* or *ftrylockfile()* functions.

16144 **RETURN VALUE**

16145        See *getc()*, *getchar()*, *putc()*, and *putchar()*.

16146 **ERRORS**

16147        No errors are defined.

16148 **EXAMPLES**

16149        None.

16150 **APPLICATION USAGE**

16151        Because they may be implemented as macros, *getc\_unlocked()* and *putc\_unlocked()* may treat  
16152        incorrectly a stream argument with side effects. In particular, *getc\_unlocked(\*f++)* and  
16153        *putc\_unlocked(\*f++)* do not necessarily work as expected. Therefore, use of these functions in  
16154        such situations should be preceded by the following statement as appropriate:

```
16155 #undef getc_unlocked
16156 #undef putc_unlocked
```

16157 **RATIONALE**

16158        Some I/O functions are typically implemented as macros for performance reasons (for example,  
16159        *putc()* and *getc()*). For safety, they need to be synchronized, but it is often too expensive to  
16160        synchronize on every character. Nevertheless, it was felt that the safety concerns were more  
16161        important; consequently, the *getc()*, *getchar()*, *putc()*, and *putchar()* functions are required to be  
16162        thread-safe. However, unlocked versions are also provided with names that clearly indicate the  
16163        unsafe nature of their operation but can be used to exploit their higher performance. These  
16164        unlocked versions can be safely used only within explicitly locked program regions, using  
16165        exported locking primitives. In particular, a sequence such as:

```
16166 flockfile(fileptr);
16167 putc_unlocked('1', fileptr);
16168 putc_unlocked('\n', fileptr);
16169 fprintf(fileptr, "Line 2\n");
16170 funlockfile(fileptr);
```

16171        is permissible, and results in the text sequence:

16172 1  
16173 Line 2  
16174 being printed without being interspersed with output from other threads.

16175 It would be wrong to have the standard names such as *getc()*, *putc()*, and so on, map to the  
16176 “faster, but unsafe” rather than the “slower, but safe” versions. In either case, you would still  
16177 want to inspect all uses of *getc()*, *putc()*, and so on, by hand when converting existing code.  
16178 Choosing the safe bindings as the default, at least, results in correct code and maintains the  
16179 “atomicity at the function” invariant. To do otherwise would introduce gratuitous  
16180 synchronization errors into converted code. Other routines that modify the *stdio* (**FILE\***)  
16181 structures or buffers are also safely synchronized.

16182 Note that there is no need for functions of the form *getc\_locked()*, *putc\_locked()*, and so on, since  
16183 this is the functionality of *getc()*, *putc()*, *et al.* It would be inappropriate to use a feature test  
16184 macro to switch a macro definition of *getc()* between *getc\_locked()* and *getc\_unlocked()*, since the  
16185 ISO C standard requires an actual function to exist, a function whose behavior could not be  
16186 changed by the feature test macro. Also, providing both the *xxx\_locked()* and *xxx\_unlocked()*  
16187 forms leads to the confusion of whether the suffix describes the behavior of the function or the  
16188 circumstances under which it should be used.

16189 Three additional routines, *flockfile()*, *ftrylockfile()*, and *funlockfile()* (which may be macros), are  
16190 provided to allow the user to delineate a sequence of I/O statements that are executed  
16191 synchronously.

16192 The *ungetc()* function is infrequently called relative to the other functions/macros so no  
16193 unlocked variation is needed.

16194 **FUTURE DIRECTIONS**  
16195 None.

16196 **SEE ALSO**  
16197 *getc()*, *getchar()*, *putc()*, *putchar()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
16198 <stdio.h>

16199 **CHANGE HISTORY**  
16200 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

16201 **Issue 6**  
16202 These functions are marked as part of the Thread-Safe Functions option.

16203 The Open Group corrigenda item U030/2 has been applied adding APPLICATION USAGE  
16204 describing how applications should be written to avoid the case when the functions are  
16205 implemented as macros.

16206 **NAME**

16207            getchar — get a byte from a stdin stream

16208 **SYNOPSIS**

16209            #include <stdio.h>

16210            int getchar(void);

16211 **DESCRIPTION**

16212 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
16213            conflict between the requirements described here and the ISO C standard is unintentional. This  
16214            volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

16215            The *getchar()* function shall be equivalent to *getc(stdin)*.

16216 **RETURN VALUE**

16217            Refer to *fgetc()*.

16218 **ERRORS**

16219            Refer to *fgetc()*.

16220 **EXAMPLES**

16221            None.

16222 **APPLICATION USAGE**

16223            If the integer value returned by *getchar()* is stored into a variable of type **char** and then  
16224            compared against the integer constant EOF, the comparison may never succeed, because sign-  
16225            extension of a variable of type **char** on widening to integer is implementation-defined.

16226 **RATIONALE**

16227            None.

16228 **FUTURE DIRECTIONS**

16229            None.

16230 **SEE ALSO**

16231            *getc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>

16232 **CHANGE HISTORY**

16233            First released in Issue 1. Derived from Issue 1 of the SVID.

16234 **Issue 4**

16235            The words “a character variable” are replaced by “a variable of type **char**”, to emphasize the fact  
16236            that this function deals in byte values.

16237            The following change is incorporated for alignment with the ISO C standard:

- 16238
  - The argument list is explicitly defined as **void**.

16239 **NAME**

16240 getcontext, setcontext — get and set current user context

16241 **SYNOPSIS**16242 XSI 

```
#include <ucontext.h>
```

16243 

```
int getcontext(ucontext_t *ucp);
```

16244 

```
int setcontext(const ucontext_t *ucp);
```

16245

16246 **DESCRIPTION**

16247 The *getcontext()* function shall initialize the structure pointed to by *ucp* to the current user  
 16248 context of the calling thread. The **ucontext\_t** type that *ucp* points to defines the user context and  
 16249 includes the contents of the calling thread's machine registers, the signal mask, and the current  
 16250 execution stack.

16251 The *setcontext()* function shall restore the user context pointed to by *ucp*. A successful call to  
 16252 *setcontext()* shall not return; program execution resumes at the point specified by the *ucp*  
 16253 argument passed to *setcontext()*. The *ucp* argument should be created either by a prior call to  
 16254 *getcontext()* or *makecontext()*, or by being passed as an argument to a signal handler. If the *ucp*  
 16255 argument was created with *getcontext()*, program execution continues as if the corresponding  
 16256 call of *getcontext()* had just returned. If the *ucp* argument was created with *makecontext()*,  
 16257 program execution continues with the function passed to *makecontext()*. When that function  
 16258 returns, the thread shall continue as if after a call to *setcontext()* with the *ucp* argument that was  
 16259 input to *makecontext()*. If the *uc\_link* member of the **ucontext\_t** structure pointed to by the *ucp*  
 16260 argument is equal to 0, then this context is the main context, and the thread shall exit when this  
 16261 context returns. The effects of passing a *ucp* argument obtained from any other source are  
 16262 unspecified.

16263 **RETURN VALUE**

16264 Upon successful completion, *setcontext()* shall not return and *getcontext()* shall return 0;  
 16265 otherwise, a value of -1 shall be returned.

16266 **ERRORS**

16267 No errors are defined.

16268 **EXAMPLES**

16269 None.

16270 **APPLICATION USAGE**

16271 When a signal handler is executed, the current user context is saved and a new context is  
 16272 created. If the thread leaves the signal handler via *longjmp()*, then it is unspecified whether the  
 16273 context at the time of the corresponding *setjmp()* call is restored and thus whether future calls to  
 16274 *getcontext()* provide an accurate representation of the current context, since the context restored  
 16275 by *longjmp()* does not necessarily contain all the information that *setcontext()* requires. Signal  
 16276 handlers should use *siglongjmp()* or *setcontext()* instead.

16277 Portable applications should not modify or access the *uc\_mcontext* member of **ucontext\_t**. A  
 16278 portable application cannot assume that context includes any process-wide static data, possibly  
 16279 including *errno*. Users manipulating contexts should take care to handle these explicitly when  
 16280 required.

16281 Use of contexts to create alternate stacks is not defined by this volume of IEEE Std. 1003.1-200x.

16282 **RATIONALE**

16283           None.

16284 **FUTURE DIRECTIONS**

16285           None.

16286 **SEE ALSO**16287           *bsd\_signal()*, *makecontext()*, *setjmp()*, *sigaction()*, *sigaltstack()*, *sigprocmask()*, *sigsetjmp()*, the Base  
16288           Definitions volume of IEEE Std. 1003.1-200x, <**ucontext.h**>16289 **CHANGE HISTORY**

16290           First released in Issue 4, Version 2.

16291 **Issue 5**

16292           Moved from X/OPEN UNIX extension to BASE.

16293           The following sentence was removed from the DESCRIPTION: “If the *ucp* argument was passed  
16294           to a signal handler, program execution continues with the program instruction following the  
16295           instruction interrupted by the signal.”

16296 **NAME**

16297         getcwd — get the path name of the current working directory

16298 **SYNOPSIS**

16299         #include <unistd.h>

16300         char \*getcwd(char \*buf, size\_t size);

16301 **DESCRIPTION**

16302         The *getcwd()* function shall place an absolute path name of the current working directory in the  
16303         array pointed to by *buf*, and return *buf*. The path name copied to the array shall contain no  
16304         components that are symbolic links. The *size* argument is the size in bytes of the character array  
16305         pointed to by the *buf* argument. If *buf* is a null pointer, the behavior of *getcwd()* is unspecified.

16306 **RETURN VALUE**

16307         Upon successful completion, *getcwd()* shall return the *buf* argument. Otherwise, *getcwd()* shall  
16308         return a null pointer and set *errno* to indicate the error. The contents of the array pointed to by  
16309         *buf* are then undefined.

16310 **ERRORS**

16311         The *getcwd()* function shall fail if:

16312         [EINVAL]         The *size* argument is 0.

16313         [ERANGE]         The size argument is greater than 0, but is smaller than the length of the path  
16314         name +1.

16315         The *getcwd()* function may fail if:

16316         [EACCES]         Read or search permission was denied for a component of the path name.

16317         [ENOMEM]         Insufficient storage space is available.

16318 **EXAMPLES**16319         **Determining the Absolute Path Name of the Current Working Directory**

16320         The following example returns a pointer to an array that holds the absolute path name of the  
16321         current working directory. The pointer is returned in the *ptr* variable, which points to the *buf*  
16322         array where the path name is stored.

16323         #include <stdlib.h>

16324         #include <unistd.h>

16325         ...

16326         long size;

16327         char \*buf;

16328         char \*ptr;

16329         size = pathconf(".", \_PC\_PATH\_MAX);

16330         if ((buf = (char \*)malloc((size\_t)size)) != NULL)

16331                 ptr = getcwd(buf, (size\_t)size);

16332         ...

16333 **APPLICATION USAGE**

16334         None.

16335 **RATIONALE**

16336 Since the maximum path name length is arbitrary unless `{PATH_MAX}` is defined, an  
 16337 application generally cannot supply a *buf* with *size* `{{PATH_MAX} + 1}`.

16338 Having `getcwd()` take no arguments and instead use the `malloc()` function to produce space for  
 16339 the returned argument was considered. The advantage is that `getcwd()` knows how big the  
 16340 working directory path name is and can allocate an appropriate amount of space. But the  
 16341 programmer would have to use the `free()` function to free the resulting object, or each use of  
 16342 `getcwd()` would further reduce the available memory. Also, `malloc()` and `free()` are used nowhere  
 16343 else in this volume of IEEE Std. 1003.1-200x. Finally, `getcwd()` is taken from the SVID where it  
 16344 has the two arguments used in this volume of IEEE Std. 1003.1-200x.

16345 The older function `getwd()` was rejected for use in this context because it had only a buffer  
 16346 argument and no *size* argument, and thus had no way to prevent overwriting the buffer, except  
 16347 to depend on the programmer to provide a large enough buffer.

16348 On some implementations, if *buf* is a null pointer, `getcwd()` may obtain *size* bytes of memory  
 16349 using `malloc()`. In this case, the pointer returned by `getcwd()` may be used as the argument in a  
 16350 subsequent call to `free()`. Invoking `getcwd()` with *buf* as a null pointer is not recommended in  
 16351 portable applications.

16352 If a program is operating in a directory where some (grand)parent directory does not permit  
 16353 reading, `getcwd()` may fail, as in most implementations it must read the directory to determine  
 16354 the name of the file. This can occur if search, but not read, permission is granted in an  
 16355 intermediate directory, or if the program is placed in that directory by some more privileged  
 16356 process (for example, login). Including the [EACCES] error condition makes the reporting of the  
 16357 error consistent and warns the application writer that `getcwd()` can fail for reasons beyond the  
 16358 control of the application writer or user. Some implementations can avoid this occurrence (for  
 16359 example, by implementing `getcwd()` using `pwd`, where `pwd` is a set-user-root process), thus the  
 16360 error was made optional.

16361 Because this volume of IEEE Std. 1003.1-200x permits the addition of other errors, this would be  
 16362 a common addition and yet one that applications could not be expected to deal with without  
 16363 this addition.

16364 **FUTURE DIRECTIONS**

16365 None.

16366 **SEE ALSO**

16367 `malloc()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<unistd.h>`

16368 **CHANGE HISTORY**

16369 First released in Issue 1. Derived from Issue 1 of the SVID.

16370 **Issue 4**

16371 The `<unistd.h>` header is added to the SYNOPSIS section.

16372 The [ENOMEM] error is marked as an extension.

16373 The words “as this functionality may be subject to withdrawal” have been deleted from the end  
 16374 of the last sentence in the APPLICATION USAGE section.

16375 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 16376 • The DESCRIPTION is changed to indicate that the effects of passing a null pointer in *buf* are  
 16377 undefined.

16378 **Issue 6**

16379 The following new requirements on POSIX implementations derive from alignment with the  
16380 Single UNIX Specification:

- 16381
- The [ENOMEM] optional error condition is added.

16382 **NAME**

16383 getdate — convert user format date and time

16384 **SYNOPSIS**

```
16385 xSI #include <time.h>
16386 struct tm *getdate(const char *string);
16387
```

16388 **DESCRIPTION**

16389 The *getdate()* function shall convert a string representation of a date or time into a broken-down  
 16390 time.

16391 The external variable or macro *getdate\_err* is used by *getdate()* to return error values.

16392 Templates are used to parse and interpret the input string. The templates are contained in a text  
 16393 file identified by the environment variable *DATEMSK*. The *DATEMSK* variable should be set to  
 16394 indicate the full path name of the file that contains the templates. The first line in the template  
 16395 that matches the input specification is used for interpretation and conversion into the internal  
 16396 time format.

16397 The following field descriptors are supported:

|       |                 |                                                                                     |
|-------|-----------------|-------------------------------------------------------------------------------------|
| 16398 | <code>%%</code> | Same as <code>' % '</code> .                                                        |
| 16399 | <code>%a</code> | Abbreviated weekday name.                                                           |
| 16400 | <code>%A</code> | Full weekday name.                                                                  |
| 16401 | <code>%b</code> | Abbreviated month name.                                                             |
| 16402 | <code>%B</code> | Full month name.                                                                    |
| 16403 | <code>%c</code> | Locale's appropriate date and time representation.                                  |
| 16404 | <code>%C</code> | Century number (00-99; leading zeros are permitted but not required).               |
| 16405 | <code>%d</code> | Day of month (01-31; the leading 0 is optional).                                    |
| 16406 | <code>%D</code> | Date as <code>%m/%d/%y</code> .                                                     |
| 16407 | <code>%e</code> | Same as <code>%d</code> .                                                           |
| 16408 | <code>%h</code> | Abbreviated month name.                                                             |
| 16409 | <code>%H</code> | Hour (00-23).                                                                       |
| 16410 | <code>%I</code> | Hour (01-12).                                                                       |
| 16411 | <code>%m</code> | Month number (01-12).                                                               |
| 16412 | <code>%M</code> | Minute (00-59).                                                                     |
| 16413 | <code>%n</code> | Same as <code>&lt;newline&gt;</code> .                                              |
| 16414 | <code>%p</code> | Locale's equivalent of either AM or PM.                                             |
| 16415 | <code>%r</code> | The locale's appropriate representation of time in AM and PM notation. In the POSIX |
| 16416 |                 | locale, this is equivalent to <code>%I:%M:%S %p</code> .                            |
| 16417 | <code>%R</code> | Time as <code>%H:%M</code> .                                                        |
| 16418 | <code>%S</code> | Seconds (00-61). Leap seconds are allowed but are not predictable through use of    |
| 16419 |                 | algorithms.                                                                         |

|       |    |                                                                                                             |
|-------|----|-------------------------------------------------------------------------------------------------------------|
| 16420 | %t | Same as <tab>.                                                                                              |
| 16421 | %T | Time as %H:%M:%S.                                                                                           |
| 16422 | %w | Weekday number (Sunday = 0-6).                                                                              |
| 16423 | %x | Locale's appropriate date representation.                                                                   |
| 16424 | %X | Locale's appropriate time representation.                                                                   |
| 16425 | %y | Year within century. When a century is not otherwise specified, values in the range                         |
| 16426 |    | 69-99 refer to years in the twentieth century (1969 to 1999 inclusive); values in the range                 |
| 16427 |    | 00-68 refer to years in the twenty-first century (2000 to 2068 inclusive).                                  |
| 16428 | %Y | Year as <i>ccyy</i> (for example, 1994).                                                                    |
| 16429 | %Z | Timezone name or no characters if no timezone exists. If the timezone supplied by %Z                        |
| 16430 |    | is not the timezone that <i>getdate()</i> expects, an invalid input specification error shall               |
| 16431 |    | result. The <i>getdate()</i> function calculates an expected timezone based on information                  |
| 16432 |    | supplied to the function (such as the hour, day, and month).                                                |
| 16433 |    | The match between the template and input specification performed by <i>getdate()</i> is case-               |
| 16434 |    | insensitive.                                                                                                |
| 16435 |    | The month and weekday names can consist of any combination of upper and lowercase letters.                  |
| 16436 |    | The process can request that the input date or time specification be in a specific language by              |
| 16437 |    | setting the <i>LC_TIME</i> category (see <i>setlocale()</i> ).                                              |
| 16438 |    | Leading 0s are not necessary for the descriptors that allow leading 0s. However, at most two                |
| 16439 |    | digits are allowed for those descriptors, including leading 0s. Extra whitespace in either the              |
| 16440 |    | template file or in <i>string</i> is ignored.                                                               |
| 16441 |    | The field descriptors %c, %x, and %X shall not be supported if they include unsupported field               |
| 16442 |    | descriptors.                                                                                                |
| 16443 |    | The following rules apply for converting the input specification into the internal format:                  |
| 16444 |    | • If %Z is being scanned, then <i>getdate()</i> initializes the broken-down time to be the current time     |
| 16445 |    | in the scanned timezone. Otherwise, it initializes the broken-down time based on the current                |
| 16446 |    | local time as if <i>localtime()</i> had been called.                                                        |
| 16447 |    | • If only the weekday is given, today is assumed if the given day is equal to the current day               |
| 16448 |    | and next week if it is less,                                                                                |
| 16449 |    | • If only the month is given, the current month is assumed if the given month is equal to the               |
| 16450 |    | current month and next year if it is less, and no year is given (the first day of month is                  |
| 16451 |    | assumed if no day is given),                                                                                |
| 16452 |    | • If no hour, minute, and second are given the current hour, minute, and second are assumed,                |
| 16453 |    | • If no date is given, today is assumed if the given hour is greater than the current hour and              |
| 16454 |    | tomorrow is assumed if it is less.                                                                          |
| 16455 |    | If a field descriptor specification in the DATEMSK file does not correspond to one of the field             |
| 16456 |    | descriptors above, the behavior is unspecified.                                                             |
| 16457 |    | The <i>getdate()</i> function need not be reentrant. A function that is not required to be reentrant is not |
| 16458 |    | required to be thread-safe.                                                                                 |

16459 **RETURN VALUE**

16460 Upon successful completion, `getdate()` shall return a pointer to a **struct tm**. Otherwise, it shall  
 16461 return a null pointer and set `getdate_err` to indicate the error.

16462 **ERRORS**

16463 The `getdate()` function shall fail in the following cases, setting `getdate_err` to the value shown in  
 16464 the list below. Any changes to `errno` are unspecified.

- 16465 1. The `DATEMSK` environment variable is null or undefined.
- 16466 2. The template file cannot be opened for reading.
- 16467 3. Failed to get file status information.
- 16468 4. The template file is not a regular file.
- 16469 5. An I/O error is encountered while reading the template file.
- 16470 6. Memory allocation failed (not enough memory available).
- 16471 7. There is no line in the template that matches the input.
- 16472 8. Invalid input specification. For example, February 31; or a time is specified that cannot be  
 16473 represented in a `time_t` (representing the time in seconds since the Epoch).

16474 **EXAMPLES**

- 16475 1. The following example shows the possible contents of a template:

```
16476 %m
16477 %A %B %d, %Y, %H:%M:%S
16478 %A
16479 %B
16480 %m/%d/%y %I %p
16481 %d,%m,%Y %H:%M
16482 at %A the %dst of %B in %Y
16483 run job at %I %p,%B %dnd
16484 %A den %d. %B %Y %H.%M Uhr
```

- 16485 2. The following are examples of valid input specifications for the template in Example 1:

```
16486 getdate("10/1/87 4 PM");
16487 getdate("Friday");
16488 getdate("Friday September 18, 1987, 10:30:30");
16489 getdate("24,9,1986 10:30");
16490 getdate("at monday the 1st of december in 1986");
16491 getdate("run job at 3 PM, december 2nd");
```

16492 If the `LC_TIME` category is set to a German locale that includes *freitag* as a weekday name  
 16493 and *oktober* as a month name, the following would be valid:

```
16494 getdate("freitag den 10. oktober 1986 10.30 Uhr");
```

- 16495 3. The following example shows how local date and time specification can be defined in the  
 16496 template:

16497  
16498  
16499  
16500  
16501  
16502

| Invocation                 | Line in Template |
|----------------------------|------------------|
| getdate("11/27/86")        | %m/%d/%y         |
| getdate("27.11.86")        | %d.%m.%y         |
| getdate("86-11-27")        | %y-%m-%d         |
| getdate("Friday 12:00:00") | %A %H:%M:%S      |

16503  
16504  
16505  
16506

4. The following examples help to illustrate the above rules assuming that the current date is Mon Sep 22 12:19:47 EDT 1986 and the *LC\_TIME* category is set to the default C locale:

16507  
16508  
16509  
16510  
16511  
16512  
16513  
16514  
16515  
16516  
16517  
16518  
16519  
16520

| Input        | Line in Template | Date                         |
|--------------|------------------|------------------------------|
| Mon          | %a               | Mon Sep 22 12:19:47 EDT 1986 |
| Sun          | %a               | Sun Sep 28 12:19:47 EDT 1986 |
| Fri          | %a               | Fri Sep 26 12:19:47 EDT 1986 |
| September    | %B               | Mon Sep 1 12:19:47 EDT 1986  |
| January      | %B               | Thu Jan 1 12:19:47 EST 1987  |
| December     | %B               | Mon Dec 1 12:19:47 EST 1986  |
| Sep Mon      | %b %a            | Mon Sep 1 12:19:47 EDT 1986  |
| Jan Fri      | %b %a            | Fri Jan 2 12:19:47 EST 1987  |
| Dec Mon      | %b %a            | Mon Dec 1 12:19:47 EST 1986  |
| Jan Wed 1989 | %b %a %Y         | Wed Jan 4 12:19:47 EST 1989  |
| Fri 9        | %a %H            | Fri Sep 26 09:00:00 EDT 1986 |
| Feb 10:30    | %b %H:%S         | Sun Feb 1 10:00:30 EST 1987  |
| 10:30        | %H:%M            | Tue Sep 23 10:30:00 EDT 1986 |
| 13:30        | %H:%M            | Mon Sep 22 13:30:00 EDT 1986 |

#### 16521 APPLICATION USAGE

16522 Although historical versions of *getdate()* did not require that **<time.h>** declare the external  
16523 variable *getdate\_err*, this volume of IEEE Std. 1003.1-200x does require it. The Open Group  
16524 encourages applications to remove declarations of *getdate\_err* and instead incorporate the  
16525 declaration by including **<time.h>**.

16526 Applications should use %Y (4-digit years) in preference to %y (2-digit years).

#### 16527 RATIONALE

16528 None.

#### 16529 FUTURE DIRECTIONS

16530 None.

#### 16531 SEE ALSO

16532 *ctime()*, *localtime()*, *setlocale()*, *strftime()*, *times()*, the Base Definitions volume of  
16533 IEEE Std. 1003.1-200x, **<time.h>**

#### 16534 CHANGE HISTORY

16535 First released in Issue 4, Version 2.

#### 16536 Issue 5

16537 Moved from X/OPEN UNIX extension to BASE.

16538 The last paragraph of the DESCRIPTION is added.

16539 The %C specifier is added, and the exact meaning of the %y specifier is clarified in the  
16540 DESCRIPTION.

16541 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

16542 The %R specifier is changed to follow historical practice.

16543 **Issue 6**

16544 The DESCRIPTION is updated to refer to “seconds since the Epoch” rather than “seconds since  
16545 00:00:00 UTC (Coordinated Universal Time), January 1 1970” for consistency with other *time* |  
16546 functions.

16547 **NAME**

16548           getegid — get the effective group ID

16549 **SYNOPSIS**

16550           #include <unistd.h>

16551           gid\_t getegid(void);

16552 **DESCRIPTION**

16553           The *getegid()* function shall return the effective group ID of the calling process.

16554 **RETURN VALUE**

16555           The *getegid()* function is always successful and no return value is reserved to indicate an error.

16556 **ERRORS**

16557           No errors are defined.

16558 **EXAMPLES**

16559           None.

16560 **APPLICATION USAGE**

16561           None.

16562 **RATIONALE**

16563           None.

16564 **FUTURE DIRECTIONS**

16565           None.

16566 **SEE ALSO**

16567           *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the Base  
16568           Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>

16569 **CHANGE HISTORY**

16570           First released in Issue 1. Derived from Issue 1 of the SVID.

16571 **Issue 4**

16572           The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
16573           XSI-conformant systems.

16574           The <unistd.h> header is added to the SYNOPSIS section.

16575           The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 16576
  - The argument list is explicitly defined as **void**.

16577 **Issue 6**

16578           In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

16579           The following new requirements on POSIX implementations derive from alignment with the  
16580           Single UNIX Specification:

- 16581
  - The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
16582           required for conforming implementations of previous POSIX specifications, it was not  
16583           required for UNIX applications.

16584 **NAME**

16585            getenv — get value of an environment variable

16586 **SYNOPSIS**

16587            #include &lt;stdlib.h&gt;

16588            char \*getenv(const char \*name);

16589 **DESCRIPTION**

16590 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
 16591 conflict between the requirements described here and the ISO C standard is unintentional. This  
 16592 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

16593        The *getenv()* function shall search the environment of the calling process (see the Base  
 16594 Definitions volume of IEEE Std. 1003.1-200x, Chapter 8, Environment Variables) for the  
 16595 environment variable *name* if it exists and return a pointer to the value of the environment  
 16596 variable. If the specified environment variable cannot be found, a null pointer shall be returned.  
 16597 The application shall ensure that it does not modify the string pointed to by the *getenv()*  
 16598 function. The string pointed to may be overwritten by a subsequent call to *getenv()*, *setenv()*,  
 16599 **XSI** *unsetenv()*, or *putenv()* but shall not be overwritten by a call to any other function in this volume  
 16600 of IEEE Std. 1003.1-200x.

16601        If the application modifies *environ* or the pointers to which it points, the behavior of *getenv()* is  
 16602 undefined.

16603 **CX**        The *getenv()* function need not be reentrant. A function that is not required to be reentrant is not  
 16604 required to be thread-safe.

16605 **RETURN VALUE**

16606        Upon successful completion, *getenv()* shall return a pointer to a string containing the *value* for  
 16607 the specified *name*. If the specified name cannot be found in the environment of the calling  
 16608 process, a null pointer shall be returned.

16609        The return value from *getenv()* may point to static data which may be overwritten by  
 16610 **XSI** subsequent calls to *getenv()*, *setenv()*, *unsetenv()*, or *putenv()*.

16611 **ERRORS**

16612        No errors are defined.

16613 **EXAMPLES**16614        **Getting the Value of an Environment Variable**16615        The following example gets the value of the *HOME* environment variable.

```
16616 #include <stdlib.h>
16617 ...
16618 const char *name = "HOME";
16619 char *value;
16620 value = getenv(name);
```

16621 **APPLICATION USAGE**

16622        None.

16623 **RATIONALE**

16624        The *clearenv()* function was considered but rejected. The *putenv()* function has now been  
 16625 included for alignment with the Single UNIX Specification.

16626 The *getenv()* function is inherently not reentrant because it returns a value pointing to static  
16627 data.

16628 Conforming applications are required not to modify *environ* directly, but to use only the  
16629 functions described here to manipulate the process environment as an abstract object. Thus, the  
16630 implementation of the environment access functions has complete control over the data  
16631 structure used to represent the environment (subject to the requirement that *environ* be  
16632 maintained as a list of strings with embedded equal signs for applications that wish to scan the  
16633 environment). This constraint allows the implementation to properly manage the memory it  
16634 allocates, either by using allocated storage for all variables (copying them on the first invocation  
16635 of *setenv()* or *unsetenv()*), or keeping track of which strings are currently in allocated space and  
16636 which are not, via a separate table or some other means. This enables the implementation to free  
16637 any allocated space used by strings (and perhaps the pointers to them) stored in *environ* when  
16638 *unsetenv()* is called. A C runtime start-up procedure (that which invokes *main()* and perhaps  
16639 initializes *environ*) can also initialize a flag indicating that none of the environment has yet been  
16640 copied to allocated storage, or that the separate table has not yet been initialized.

16641 In fact, for higher performance of *getenv()*, the implementation could also maintain a separate  
16642 copy of the environment in a data structure that could be searched much more quickly (such as  
16643 an indexed hash table, or a binary tree), and update both it and the linear list at *environ* when  
16644 *setenv()* or *unsetenv()* is invoked.

16645 Performance of *getenv()* can be important for applications which have large numbers of  
16646 environment variables. Typically, applications like this use the environment as a resource  
16647 database of user-configurable parameters. The fact that these variables are in the user's shell  
16648 environment usually means that any other program that uses environment variables (such as *ls*,  
16649 which attempts to use *COLUMNS*, or really almost any utility (*LANG*, *LC\_ALL*, and so on) is  
16650 similarly slowed down by the linear search through the variables.

16651 An implementation that maintains separate data structures, or even one that manages the  
16652 memory it consumes, is not currently required as it was thought it would reduce consensus  
16653 among implementors who do not want to change their historical implementations.

16654 The POSIX Threads Extension states that multi-threaded applications must not modify *environ*  
16655 directly, and that IEEE Std. 1003.1-200x is providing functions which such applications can use  
16656 in the future to manipulate the environment in a thread-safe manner. Thus, moving away from  
16657 application use of *environ* is desirable from that standpoint as well.

#### 16658 FUTURE DIRECTIONS

16659 None.

#### 16660 SEE ALSO

16661 *exec*, *putenv()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>`, the Base  
16662 Definitions volume of IEEE Std. 1003.1-200x, Chapter 8, Environment Variables

#### 16663 CHANGE HISTORY

16664 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 16665 Issue 4

16666 The DESCRIPTION is updated to indicate that the return string must not be modified by an  
16667 application, may be overwritten by subsequent calls to *getenv()* or *putenv()*, and is not  
16668 overwritten by calls to other XSI system interfaces.

16669 A reference to *putenv()* has also been added to the APPLICATION USAGE section.

16670 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 16671
  - The type of argument *name* is changed from **char\*** to **const char\***.
- 16672 **Issue 5**
- 16673 Normative text previously in the APPLICATION USAGE section is moved to the RETURN
- 16674 VALUE section.
- 16675 A note indicating that this function need not be reentrant is added to the DESCRIPTION.
- 16676 **Issue 6**
- 16677 The following changes were made to align with the IEEE P1003.1a draft standard:
  - References added to the new *setenv()* and *unsetenv()* functions.
- 16678
- 16679 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

16680 **NAME**

16681           geteuid — get the effective user ID

16682 **SYNOPSIS**

16683           #include <unistd.h>

16684           uid\_t geteuid(void);

16685 **DESCRIPTION**

16686           The *geteuid()* function shall return the effective user ID of the calling process.

16687 **RETURN VALUE**

16688           The *geteuid()* function is always successful and no return value is reserved to indicate an error.

16689 **ERRORS**

16690           No errors are defined.

16691 **EXAMPLES**

16692           None.

16693 **APPLICATION USAGE**

16694           None.

16695 **RATIONALE**

16696           None.

16697 **FUTURE DIRECTIONS**

16698           None.

16699 **SEE ALSO**

16700           *getegid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the Base  
16701           Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>

16702 **CHANGE HISTORY**

16703           First released in Issue 1. Derived from Issue 1 of the SVID.

16704 **Issue 4**

16705           The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
16706           XSI-conformant systems.

16707           The <unistd.h> header is added to the SYNOPSIS section.

16708           The following change is incorporated for alignment with the ISO POSIX-1 standard:

16709           

- The argument list is explicitly defined as **void**.

16710 **Issue 6**

16711           In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

16712           The following new requirements on POSIX implementations derive from alignment with the  
16713           Single UNIX Specification:

16714           

- The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
16715           required for conforming implementations of previous POSIX specifications, it was not  
16716           required for UNIX applications.

16717 **NAME**

16718           getgid — get the real group ID

16719 **SYNOPSIS**

16720           #include <unistd.h>

16721           gid\_t getgid(void);

16722 **DESCRIPTION**

16723           The *getgid()* function shall return the real group ID of the calling process.

16724 **RETURN VALUE**

16725           The *getgid()* function is always successful and no return value is reserved to indicate an error.

16726 **ERRORS**

16727           No errors are defined.

16728 **EXAMPLES**

16729           None.

16730 **APPLICATION USAGE**

16731           None.

16732 **RATIONALE**

16733           None.

16734 **FUTURE DIRECTIONS**

16735           None.

16736 **SEE ALSO**

16737           *getegid()*, *geteuid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the Base  
16738           Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>

16739 **CHANGE HISTORY**

16740           First released in Issue 1. Derived from Issue 1 of the SVID.

16741 **Issue 4**

16742           The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
16743           XSI-conformant systems.

16744           The <unistd.h> header is added to the SYNOPSIS section.

16745           The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 16746
  - The argument list is explicitly defined as **void**.

16747 **Issue 6**

16748           In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

16749           The following new requirements on POSIX implementations derive from alignment with the  
16750           Single UNIX Specification:

- 16751
  - The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
16752           required for conforming implementations of previous POSIX specifications, it was not  
16753           required for UNIX applications.

16754 **NAME**

16755           getgrent — get the group database entry

16756 **SYNOPSIS**

16757 xSI       #include <grp.h>

16758           struct group \*getgrent(void);

16759

16760 **DESCRIPTION**

16761           Refer to *endgrent()*.

16762 **NAME**

16763 getgrgid, getgrgid\_r — get group database entry for a group ID

16764 **SYNOPSIS**

16765 #include &lt;grp.h&gt;

16766 struct group \*getgrgid(gid\_t gid);

16767 TSF int getgrgid\_r(gid\_t gid, struct group \*grp, char \*buffer,

16768 size\_t bufsize, struct group \*\*result);

16769

16770 **DESCRIPTION**16771 The *getgrgid()* function shall search the group database for an entry with a matching *gid*.16772 The *getgrgid()* function need not be reentrant. A function that is not required to be reentrant is  
16773 not required to be thread-safe.16774 TSF The *getgrgid\_r()* function updates the **group** structure pointed to by *grp* and stores a pointer to  
16775 that structure at the location pointed to by *result*. The structure contains an entry from the  
16776 group database with a matching *gid*. Storage referenced by the group structure is allocated from  
16777 the memory provided with the *buffer* parameter, which is *bufsize* characters in size. The  
16778 maximum size needed for this buffer can be determined with the `{_SC_GETGR_R_SIZE_MAX}`  
16779 *sysconf()* parameter. A NULL pointer is returned at the location pointed to by *result* on error or if  
16780 the requested entry is not found.16781 **RETURN VALUE**16782 Upon successful completion, *getgrgid()* shall return a pointer to a **struct group** with the structure  
16783 defined in `<grp.h>` with a matching entry if one is found. The *getgrgid()* function shall return a  
16784 null pointer if either the requested entry was not found, or an error occurred. On error, *errno*  
16785 shall be set to indicate the error.16786 The return value may point to a static area which is overwritten by a subsequent call to  
16787 *getgrent()*, *getgrgid()*, or *getgrnam()*.16788 TSF If successful, the *getgrgid\_r()* function shall return zero; otherwise, an error number shall be  
16789 returned to indicate the error.16790 **ERRORS**16791 The *getgrgid()* and *getgrgid\_r()* functions may fail if:

16792 [EIO] An I/O error has occurred.

16793 [EINTR] A signal was caught during *getgrgid()*.16794 [EMFILE] `{OPEN_MAX}` file descriptors are currently open in the calling process.

16795 [ENFILE] The maximum allowable number of files is currently open in the system.

16796 TSF The *getgrgid\_r()* function may fail if:16797 TSF [ERANGE] Insufficient storage was supplied via *buffer* and *bufsize* to contain the data to  
16798 be referenced by the resulting **group** structure.

16799 **EXAMPLES**16800 **Finding an Entry in the Group Database**

16801 The following example uses *getgrgid()* to search the group database for a group ID that was  
 16802 previously stored in a **stat** structure, then prints out the group name if it is found. If the group is  
 16803 not found, the program prints the numeric value of the group for the entry.

```
16804 #include <sys/types.h>
16805 #include <grp.h>
16806 #include <stdio.h>
16807 ...
16808 struct stat statbuf;
16809 struct group *grp;
16810 ...
16811 if ((grp = getgrgid(statbuf.st_gid)) != NULL)
16812 printf(" %-8.8s", grp->gr_name);
16813 else
16814 printf(" %-8d", statbuf.st_gid);
16815 ...
```

16816 **APPLICATION USAGE**

16817 Applications wishing to check for error situations should set *errno* to 0 before calling *getgrgid()*.  
 16818 If *errno* is set on return, an error occurred.

16819 The *getgrgid\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
 16820 of possibly using a static data area that may be overwritten by each call.

16821 **RATIONALE**

16822 None.

16823 **FUTURE DIRECTIONS**

16824 None.

16825 **SEE ALSO**

16826 *endgrent()*, *getgrnam()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<grp.h>**,  
 16827 **<limits.h>**, **<sys/types.h>**

16828 **CHANGE HISTORY**

16829 First released in Issue 1. Derived from System V Release 2.0.

16830 **Issue 4**

16831 The DESCRIPTION is clarified.

16832 In the RETURN VALUE section, the reference to the setting of *errno* is marked as an extension.

16833 The errors [EIO], [EINTR], [EMFILE], and [ENFILE] are marked as extensions.

16834 A note is added to the APPLICATION USAGE section advising how applications should check  
 16835 for errors.

16836 The **<sys/types.h>** header is added as optional (OH); this header need not be included on XSI-  
 16837 conformant systems.

16838 **Issue 5**

16839 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
 16840 VALUE section.

16841 The *getgrgid\_r()* function is included for alignment with the POSIX Threads Extension.

- 16842 A note indicating that the *getgrgid()* function need not be reentrant is added to the  
16843 DESCRIPTION.
- 16844 **Issue 6**
- 16845 The *getgrgid\_r()* function is marked as part of the Thread-Safe Functions option.
- 16846 The Open Group corrigenda item U028/3 has been applied correcting text in the DESCRIPTION  
16847 describing matching the *gid*.
- 16848 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.
- 16849 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.
- 16850 The following new requirements on POSIX implementations derive from alignment with the  
16851 Single UNIX Specification:
- 16852 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
16853 required for conforming implementations of previous POSIX specifications, it was not  
16854 required for UNIX applications.
  - 16855 • In the RETURN VALUE section, the requirement to set *errno* on error is added.
  - 16856 • The [EIO], [EINTR], [EMFILE], and [ENFILE] optional error conditions are added.
- 16857 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
16858 its avoidance of possibly using a static data area.

## 16859 NAME

16860 getgrnam, getgrnam\_r — search group database for a name

## 16861 SYNOPSIS

16862 #include <grp.h>

16863 struct group \*getgrnam(const char \*name);

16864 TSF int getgrnam\_r(const char \*name, struct group \*grp, char \*buffer,

16865 size\_t bufsize, struct group \*\*result);

16866

## 16867 DESCRIPTION

16868 The *getgrnam()* function shall search the group database for an entry with a matching *name*.

16869 The *getgrnam()* function need not be reentrant. A function that is not required to be reentrant is  
16870 not required to be thread-safe.

16871 TSF The *getgrnam\_r()* function updates the **group** structure pointed to by *grp* and stores a pointer to  
16872 that structure at the location pointed to by *result*. The structure contains an entry from the  
16873 group database with a matching *gid* or *name*. Storage referenced by the group structure is  
16874 allocated from the memory provided with the *buffer* parameter, which is *bufsize* characters in  
16875 size. The maximum size needed for this buffer can be determined with the  
16876 `{_SC_GETGR_R_SIZE_MAX} sysconf()` parameter. A NULL pointer is returned at the location  
16877 pointed to by *result* on error or if the requested entry is not found.

## 16878 RETURN VALUE

16879 The *getgrnam()* function shall return a pointer to a **struct group** with the structure defined in  
16880 <grp.h> with a matching entry if one is found. The *getgrnam()* function shall return a null  
16881 pointer if either the requested entry was not found, or an error occurred. On error, *errno* shall be  
16882 set to indicate the error.

16883 The return value may point to a static area which is overwritten by a subsequent call to  
16884 *getgrent()*, *getgrgid()*, or *getgrnam()*.

16885 TSF If successful, the *getgrnam\_r()* function shall return zero; otherwise, an error number shall be  
16886 returned to indicate the error.

## 16887 ERRORS

16888 The *getgrnam()* and *getgrnam\_r()* functions may fail if:

16889 [EIO] An I/O error has occurred.

16890 [EINTR] A signal was caught during *getgrnam()*.

16891 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

16892 [ENFILE] The maximum allowable number of files is currently open in the system.

16893 The *getgrnam\_r()* function may fail if:

16894 TSF [ERANGE] Insufficient storage was supplied via *buffer* and *bufsize* to contain the data to  
16895 be referenced by the resulting **group** structure.

16896 **EXAMPLES**

16897       None.

16898 **APPLICATION USAGE**

16899       Applications wishing to check for error situations should set *errno* to 0 before calling *getgrnam()*.  
 16900       If *errno* is set on return, an error occurred.

16901       The *getgrnam\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
 16902       of possibly using a static data area that may be overwritten by each call.

16903 **RATIONALE**

16904       None.

16905 **FUTURE DIRECTIONS**

16906       None.

16907 **SEE ALSO**

16908       *endgrent()*, *getgrgid()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<grp.h>`, `<limits.h>`,  
 16909       `<sys/types.h>`

16910 **CHANGE HISTORY**

16911       First released in Issue 1. Derived from System V Release 2.0.

16912 **Issue 4**

16913       The DESCRIPTION is clarified.

16914       The `<sys/types.h>` header is added as optional (OH); this header need not be included on XSI-  
 16915       conformant systems.

16916       In the RETURN VALUE section, reference to the setting of *errno* is marked as an extension.

16917       The errors [EIO], [EINTR], [EMFILE], and [ENFILE] are marked as extensions.

16918       A note is added to the APPLICATION USAGE section advising how applications should check  
 16919       for errors.

16920       The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 16921       • The type of argument *name* is changed from `char*` to `const char*`.

16922 **Issue 5**

16923       Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
 16924       VALUE section.

16925       The *getgrnam\_r()* function is included for alignment with the POSIX Threads Extension.

16926       A note indicating that the *getgrnam()* function need not be reentrant is added to the  
 16927       DESCRIPTION.

16928 **Issue 6**16929       The *getgrnam\_r()* function is marked as part of the Thread-Safe Functions option.

16930       In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

16931       In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

16932       The following new requirements on POSIX implementations derive from alignment with the  
 16933       Single UNIX Specification:

- 16934       • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
 16935       required for conforming implementations of previous POSIX specifications, it was not  
 16936       required for UNIX applications.

- 16937           • In the RETURN VALUE section, the requirement to set *errno* on error is added.
- 16938           • The [EIO], [EINTR], [EMFILE], and [ENFILE] optional error conditions are added.
- 16939           The APPLICATION USAGE section is updated to include a note on the thread-safe function and
- 16940           its avoidance of possibly using a static data area.

16941 **NAME**

16942           getgroups — get supplementary group IDs

16943 **SYNOPSIS**

16944           #include &lt;unistd.h&gt;

16945           int getgroups(int *gidsetsize*, gid\_t *grouplist*[]);16946 **DESCRIPTION**

16947       The *getgroups()* function fills in the array *grouplist* with the current supplementary group IDs of the calling process. It is implementation-defined whether *getgroups()* also returns the effective group ID in the *grouplist* array.

16950       The *gidsetsize* argument specifies the number of elements in the array *grouplist*. The actual number of group IDs stored in the array is returned. The values of array entries with indices greater than or equal to the value returned are undefined.

16953       If *gidsetsize* is 0, *getgroups()* shall return the number of group IDs that it would otherwise return without modifying the array pointed to by *grouplist*.

16955       If the effective group ID of the process is returned with the supplementary group IDs, the value returned shall always be greater than or equal to one and less than or equal to the value of {NGROUPS\_MAX}+1.

16958 **RETURN VALUE**

16959       Upon successful completion, the number of supplementary group IDs shall be returned. A return value of -1 indicates failure and *errno* shall be set to indicate the error.

16961 **ERRORS**16962       The *getgroups()* function shall fail if:

|       |          |                                                                                  |
|-------|----------|----------------------------------------------------------------------------------|
| 16963 | [EINVAL] | The <i>gidsetsize</i> argument is non-zero and less than the number of group IDs |
| 16964 |          | that would have been returned.                                                   |

16965 **EXAMPLES**16966       **Getting the Supplementary Group IDs of the Calling Process**

16967       The following example places the current supplementary group IDs of the calling process into the *group* array.

```
16969 #include <sys/types.h>
16970 #include <unistd.h>
16971 ...
16972 gid_t *group;
16973 int nogroups;
16974 long ngroups_max;

16975 ngroups_max = sysconf(_SC_NGROUPS_MAX);
16976 group = (gid_t *)malloc(ngroups_max *sizeof(gid_t));
16977 ngroups = getgroups(ngroups_max, group);
```

16978 **APPLICATION USAGE**

16979       None.

16980 **RATIONALE**

16981       The related function *setgroups()* is a privileged operation and therefore is not covered by this volume of IEEE Std. 1003.1-200x.

16983 As implied by the definition of supplementary groups, the effective group ID may appear in the  
16984 array returned by *getgroups()* or it may be returned only by *getegid()*. Duplication may exist, but  
16985 the application needs to call *getegid()* to be sure of getting all of the information. Various  
16986 implementation variations and administrative sequences cause the set of groups appearing in  
16987 the result of *getgroups()* to vary in order and as to whether the effective group ID is included,  
16988 even when the set of groups is the same (in the mathematical sense of “set”). (The history of a  
16989 process and its parents could affect the details of result.)

16990 Applications writers should note that {NGROUPS\_MAX} is not necessarily a constant on all  
16991 implementations.

#### 16992 FUTURE DIRECTIONS

16993 None.

#### 16994 SEE ALSO

16995 *getegid()*, *setgid()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, |  
16996 <unistd.h>

#### 16997 CHANGE HISTORY

16998 First released in Issue 3.

16999 Entry included for alignment with the POSIX.1-1988 standard.

#### 17000 Issue 4

17001 The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
17002 XSI-conformant systems.

17003 The <unistd.h> header is added to the SYNOPSIS section.

17004 The following change is incorporated for alignment with the FIPS requirements:

- 17005 • A return value of 0 is no longer permitted, because {NGROUPS\_MAX} cannot be 0.

#### 17006 Issue 5

17007 Normative text previously in the APPLICATION USAGE section is moved to the  
17008 DESCRIPTION.

#### 17009 Issue 6

17010 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

17011 The following new requirements on POSIX implementations derive from alignment with the  
17012 Single UNIX Specification:

- 17013 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
17014 required for conforming implementations of previous POSIX specifications, it was not  
17015 required for UNIX applications.

- 17016 • A return value of 0 is not permitted, because {NGROUPS\_MAX} cannot be 0. This is a FIPS  
17017 requirement.

17018 The following changes were made to align with the IEEE P1003.1a draft standard:

- 17019 • Explanation added that the effective group ID may be included in the supplementary group  
17020 list.

## 17021 NAME

17022 gethostbyaddr (LEGACY), gethostbyname (LEGACY), getipnodebyaddr, getipnodebyname, —  
 17023 network host database functions

## 17024 SYNOPSIS

```
17025 #include <netdb.h>

17026 struct hostent *gethostbyaddr(const void *addr, socklen_t len,
17027 int type);
17028 struct hostent *gethostbyname(const char *name);
17029 struct hostent *getipnodebyaddr(const void *restrict addr, socklen_t len, |
17030 int type, int *restrict error_num);
17031 struct hostent *getipnodebyname(const char *name, int type, int flags, |
17032 int *error_num);
```

## 17033 DESCRIPTION

17034 These functions enable applications to retrieve information about hosts. This information is  
 17035 considered to be stored in a database that can be accessed sequentially or randomly.  
 17036 Implementation of this database is unspecified.

17037 **Note:** In many cases it is implemented by the Domain Name System, as documented in  
 17038 RFC 1034, RFC 1035, and RFC 1886.

17039 Entries are returned in **hostent** structures.

17040 The *gethostbyaddr()* function shall return an entry containing addresses of address family *type* for  
 17041 the host with address *addr*. *len* contains the length of the address pointed to by *addr*. The  
 17042 *gethostbyaddr()* function need not be reentrant. A function that is not required to be reentrant is  
 17043 not required to be thread-safe.

17044 The *gethostbyname()* function shall return an entry containing addresses of address family  
 17045 AF\_INET for the host with name *name*. The *gethostbyname()* function need not be reentrant. A  
 17046 function that is not required to be reentrant is not required to be thread-safe.

17047 The *getipnodebyaddr()* function shall return the entry containing addresses of address family *type*  
 17048 for the host with address *addr*, opening a connection to the database if necessary. The *len*  
 17049 argument contains the length of the address pointed to by *addr*. If an error occurs, the  
 17050 appropriate error code is returned in *error\_num*. The *getipnodebyaddr()* function is thread-safe.

17051 The *getipnodebyname()* function shall return the entry containing addresses of address family  
 17052 *type* for the host with name *name*, opening a connection to the database if necessary. The *flags*  
 17053 argument affects what information is returned. If an error occurs, the appropriate error code is  
 17054 returned in *error\_num*. The *getipnodebyname()* function is thread-safe.

17055 The *addr* argument of *gethostbyaddr()* or *getipnodebyaddr()* shall be an **in\_addr** structure when  
 17056 IP6 *type* is AF\_INET, and shall be an **in6\_addr** structure when *type* is AF\_INET6. It contains a binary  
 17057 format (that is, not null-terminated) address in network byte order. The *gethostbyaddr()* function  
 17058 is not guaranteed to return addresses of address families other than AF\_INET, even when such  
 17059 addresses exist in the database.

17060 If *gethostbyaddr()* or *getipnodebyaddr()* returns a record, then its *h\_addrtype* field is the same as the  
 17061 *type* argument that was passed to the function, and its *h\_addr\_list* field lists a single address that  
 17062 IP6 is a copy of the *addr* argument that was passed to the function. If *type* is AF\_INET6 and *addr* is  
 17063 an IPv4-mapped IPv6 address or an IPv4-compatible IPv6 address, then the *h\_name* and *h\_aliases*  
 17064 fields are those that would have been returned for address family AF\_INET and address equal to  
 17065 the last four bytes of *addr*.

17066 If *gethostbyaddr()* or *getipnodebyaddr()* are called with *addr* containing the IPv6 unspecified  
 17067 address (all bytes zero), then no query is performed and the function fails with  
 17068 [HOST\_NOT\_FOUND].

17069 The *name* argument of *getipnodebyname()* shall be either a node name or a numeric address  
 17070 string. For IPv4, a numeric address string shall be in the dotted-decimal notation described in  
 17071 IP6 *inet\_addr()*. For IPv6, a numeric address string shall be in one of the standard IPv6 text forms  
 17072 described in *inet\_ntop()*. The *name* argument of *gethostbyname()* shall be a node name; the  
 17073 behavior of *gethostbyname()* when passed a numeric address string is unspecified.

17074 IP6 If *name* is a dotted-decimal IPv4 address and *af* equals AF\_INET, or *name* is an IPv6 hex address  
 17075 and *af* equals AF\_INET6, the members of the returned **hostent** structure are as follows:

17076 *h\_name* Points to a copy of the *name* argument.

17077 *h\_aliases* A NULL pointer.

17078 *h\_addrtype* A copy of the *type* argument.

17079 IP6 *h\_length* Either 4 (for AF\_INET) or 16 (for AF\_INET6).

17080 IP6 *h\_addr\_list*[0] A pointer to the 4-byte or 16-byte binary address.

17081 *h\_addr\_list*[1] A NULL pointer.

17082 IP6 If *name* is a dotted-decimal IPv4 address and *af* equals AF\_INET6 and AI\_V4MAPPED is set in  
 17083 *flags*, an IPv4-mapped IPv6 address is returned, and the members are as follows:

17084 *h\_name* Points to an IPv6 hex address containing the IPv4-mapped IPv6 address.

17085 *h\_aliases* A NULL pointer.

17086 *h\_addrtype* AF\_INET6.

17087 *h\_length* 16.

17088 *h\_addr\_list*[0] A pointer to the 16-byte binary address.

17089 *h\_addr\_list*[1] A NULL pointer.

17090 If *name* is a dotted-decimal IPv4 address and *af* equals AF\_INET6 and AI\_V4MAPPED is not set,  
 17091 then NULL is returned with [HOST\_NOT\_FOUND].

17092 It is an error when *name* is an IPv6 hex address and *af* equals AF\_INET. The function's return  
 17093 value is a NULL pointer with the [HOST\_NOT\_FOUND] error.

17094 If *name* is not a numeric address string and is an alias for a valid host name, then *gethostbyname()*  
 17095 or *getipnodebyname()* return information about the host name to which the alias refers, and *name*  
 17096 is included in the list of aliases returned.

17097 If *name* is a node name, then operation of the *getipnodebyname()* function is modified by the value  
 17098 of the *flags* argument, as follows:

17099 • If *flags* is 0 and *type* is AF\_INET, then a query is made for IPv4 addresses. If it is successful,  
 17100 the IPv4 addresses are returned and the *h\_length* member of the **hostent** structure shall have  
 17101 IP6 a value of 4. Otherwise, the function shall return a NULL pointer.

17102 • If *flags* is 0 and if *type* is AF\_INET6, then a query is made for IPv6 addresses. If it is successful,  
 17103 the IPv6 addresses are returned and the *h\_length* member of the **hostent** structure shall have  
 17104 a value of 16. If unsuccessful, the function shall return a NULL pointer.

17105 • If the AI\_V4MAPPED flag is set and *type* is AF\_INET6, then a query is made for IPv6  
 17106 addresses. If it is successful, the IPv6 addresses are returned, and no query is made for IPv4

17107 addresses. If it is not successful, a query is made for IPv4 addresses and any found are  
 17108 returned as IPv4-mapped IPv6 addresses. *h\_length* shall have a value of 16 in either case of  
 17109 addresses being returned. The AI\_V4MAPPED flag is ignored unless *type* is AF\_INET6.

- 17110 • If the AI\_ALL and AI\_V4MAPPED flags are both set and *type* is AF\_INET6, then a query is  
 17111 made for IPv6 addresses, and any found are returned. Another query is then made for IPv4  
 17112 addresses, and any found are returned as IPv4-mapped IPv6 addresses, and *h\_length* is 16.  
 17113 Only if both queries fail does the function return a NULL pointer. This flag is ignored unless  
 17114 *type* is AF\_INET6.
- 17115 • The AI\_ADDRCONFIG flag specifies that a query for IPv6 addresses should be made only if  
 17116 the node has at least one IPv6 source address configured, and that a query for IPv4 addresses  
 17117 should be made only if the node has at least one IPv4 source address configured.
- 17118 • If the AI\_V4MAPPED and AI\_ADDRCONFIG flags are both set and *type* is AF\_INET6, then:
  - 17119 — If the node has at least one IPv6 source address configured, a query is made for IPv6  
 17120 addresses.
  - 17121 — If it is successful, the IPv6 addresses are returned and no query is made for IPv4  
 17122 addresses.
  - 17123 — If the node has no IPv6 source address configured, or if the query for IPv6 addresses is not  
 17124 successful, then if the node has at least one IPv4 source address configured, a query is  
 17125 made for IPv4 addresses and any found are returned as IPv4-mapped IPv6 addresses.

17126 *h\_length* shall have a value of 16 in either case of addresses being returned.

- 17127 • Macro AI\_DEFAULT is defined as the logical OR of AI\_V4MAPPED and AI\_ADDRCONFIG.

17128 **Note:** It is intended that setting *flags* to AI\_DEFAULT be appropriate for most  
 17129 applications.

17130 **RETURN VALUE**

17131 Upon successful completion, these functions shall return a pointer to a **hostent** structure if the  
 17132 requested entry was found, and a null pointer if the end of the database was reached or the  
 17133 requested entry was not found.

17134 Upon unsuccessful completion, *getipnodebyaddr()* and *getipnodebyname()* shall set their *error\_num*  
 17135 argument to indicate the error, while *gethostbyaddr()* and *gethostbyname()* shall set *h\_errno* to  
 17136 indicate it.

17137 **ERRORS**

17138 These functions shall fail in the following cases. The *getipnodebyaddr()* and *getipnodebyname()*  
 17139 functions shall return the value shown in the list below in *error\_num*; the *gethostbyaddr()* and  
 17140 *gethostbyname()* functions shall set *h\_errno* to that value. Any changes to *errno* are unspecified.

17141 [HOST\_NOT\_FOUND]  
 17142 No such host is known.

17143 [NO\_DATA] The server recognized the request and the name, but no address is available.  
 17144 Another type of request to the name server for the domain might return an  
 17145 answer.

17146 [NO\_RECOVERY]  
 17147 An unexpected server failure occurred which cannot be recovered.

17148 [TRY\_AGAIN] A temporary and possibly transient error occurred, such as a failure of a  
 17149 server to respond.

17150 **EXAMPLES**

17151           None.

17152 **APPLICATION USAGE**

17153           The **hostent** structure returned by *getipnodebyaddr()* and *getipnodebyname()*, and any structures  
17154           pointed to from those structures, are dynamically allocated. Applications should call  
17155           *freehostent()* to free the memory used by these structures.

17156           The *gethostbyaddr()*, and *gethostbyname()* functions may return pointers to static data, which may  
17157           be overwritten by subsequent calls to any of these functions. Applications shall not call  
17158           *freehostent()* for this area.

17159           The *getipnodebyaddr()* function is preferred over the *gethostbyaddr()* function.

17160           The *getipnodebyname()* function is preferred over the *gethostbyname()* function.

17161 **RATIONALE**

17162           None.

17163 **FUTURE DIRECTIONS**

17164           The *gethostbyaddr()* and *gethostbyname()* functions may be withdrawn in a future version.

17165 **SEE ALSO**

17166           *endhostent()*, *freehostent()*, *endservent()*, *inet\_addr()*, the Base Definitions volume of  
17167           IEEE Std. 1003.1-200x, <netdb.h>

17168 **CHANGE HISTORY**

17169           First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

17170           The **restrict** keyword is added to the *getipnodebyaddr()* prototype for alignment with the  
17171           ISO/IEC 9899:1999 standard.

17172 **NAME**

17173       gethostbyname — network host database functions

17174 **SYNOPSIS**

17175       #include <netdb.h>

17176       struct hostent \*gethostbyname(const char \*name);

17177 **DESCRIPTION**

17178       Refer to *gethostbyaddr()*.

17179 **NAME**

17180       gethostent — network host database functions

17181 **SYNOPSIS**

17182       #include <netdb.h>

17183       struct hostent \*gethostent(void);

17184 **DESCRIPTION**

17185       Refer to *endhostent()*.

17186 **NAME**

17187           gethostid — get an identifier for the current host

17188 **SYNOPSIS**

17189 XSI       #include <unistd.h>

17190           long gethostid(void);

17191

17192 **DESCRIPTION**

17193           The *gethostid()* function retrieves a 32-bit identifier for the current host.

17194 **RETURN VALUE**

17195           Upon successful completion, *gethostid()* shall return an identifier for the current host.

17196 **ERRORS**

17197           No errors are defined.

17198 **EXAMPLES**

17199           None.

17200 **APPLICATION USAGE**

17201           This volume of IEEE Std. 1003.1-200x does not define the domain in which the return value is  
17202           unique.

17203 **RATIONALE**

17204           None.

17205 **FUTURE DIRECTIONS**

17206           None.

17207 **SEE ALSO**

17208           *random()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <unistd.h>

17209 **CHANGE HISTORY**

17210           First released in Issue 4, Version 2.

17211 **Issue 5**

17212           Moved from X/OPEN UNIX extension to BASE.

17213 **NAME**

17214       gethostname — get name of current host

17215 **SYNOPSIS**

17216       #include <unistd.h>

17217       int gethostname(char \*name, socklen\_t namelen);

17218 **DESCRIPTION**

17219       The *gethostname()* function shall return the standard host name for the current machine. The  
17220       *namelen* argument shall specify the size of the array pointed to by the *name* argument. The  
17221       returned name shall be null-terminated, except that if *namelen* is an insufficient length to hold  
17222       the host name, then the returned name shall be truncated and it is unspecified whether the  
17223       returned name is null-terminated.

17224       Host names are limited to 255 bytes.

17225 **RETURN VALUE**

17226       Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned.

17227 **ERRORS**

17228       No errors are defined.

17229 **EXAMPLES**

17230       None.

17231 **APPLICATION USAGE**

17232       None.

17233 **RATIONALE**

17234       None.

17235 **FUTURE DIRECTIONS**

17236       None.

17237 **SEE ALSO**

17238       *gethostid()*, *uname()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <unistd.h>

17239 **CHANGE HISTORY**

17240       First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

17241 **NAME**

17242       getipnodebyaddr — network host database functions

17243 **SYNOPSIS**

17244       #include &lt;netdb.h&gt;

17245       struct hostent \*getipnodebyaddr(const void \*restrict *addr*, socklen\_t *len*, |  
17246                   int *type*, int \*restrict *error\_num*); |17247 **DESCRIPTION**17248       Refer to *gethostbyaddr()*. |

17249 **NAME**

17250       getipnodebyname — network host database functions

17251 **SYNOPSIS**

17252       #include <netdb.h>

17253       struct hostent \*getipnodebyname(const char \*name, int type, int flags,  
17254                                   int \*error\_num);

17255 **DESCRIPTION**

17256       Refer to *gethostbyaddr()*.

17257 **NAME**

17258 getitimer, setitimer — get or set value of interval timer

17259 **SYNOPSIS**

```
17260 XSI #include <sys/time.h>
17261
17261 int getitimer(int which, struct itimerval *value);
17262 int setitimer(int which, const struct itimerval *restrict value,
17263 struct itimerval *restrict ovalue);
17264
```

17265 **DESCRIPTION**

17266 The *getitimer()* function shall store the current value of the timer specified by *which* into the  
 17267 structure pointed to by *value*. The *setitimer()* function shall set the timer specified by *which* to  
 17268 the value specified in the structure pointed to by *value*, and if *ovalue* is not a null pointer, stores  
 17269 the previous value of the timer in the structure pointed to by *ovalue*.

17270 A timer value is defined by the **itimerval** structure, specified in `<sys/time.h>`. If *it\_value* is non-  
 17271 zero, it shall indicate the time to the next timer expiration. If *it\_interval* is non-zero, it shall  
 17272 specify a value to be used in reloading *it\_value* when the timer expires. Setting *it\_value* to 0 shall  
 17273 disable a timer, regardless of the value of *it\_interval*. Setting *it\_interval* to 0 shall disable a timer  
 17274 after its next expiration (assuming *it\_value* is non-zero).

17275 Implementations may place limitations on the granularity of timer values. For each interval  
 17276 timer, if the requested timer value requires a finer granularity than the implementation supports,  
 17277 the actual timer value shall be rounded up to the next supported value.

17278 An XSI-conforming implementation provides each process with at least three interval timers,  
 17279 which are indicated by the *which* argument:

|       |                |                                                                                                                                                                                                                                                                                        |
|-------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 17280 | ITIMER_REAL    | Decrements in real time. A SIGALRM signal is delivered when this timer expires.                                                                                                                                                                                                        |
| 17281 |                |                                                                                                                                                                                                                                                                                        |
| 17282 | ITIMER_VIRTUAL | Decrements in process virtual time. It runs only when the process is executing. A SIGVTALRM signal is delivered when it expires.                                                                                                                                                       |
| 17283 |                |                                                                                                                                                                                                                                                                                        |
| 17284 | ITIMER_PROF    | Decrements both in process virtual time and when the system is running on behalf of the process. It is designed to be used by interpreters in statistically profiling the execution of interpreted programs. Each time the ITIMER_PROF timer expires, the SIGPROF signal is delivered. |
| 17285 |                |                                                                                                                                                                                                                                                                                        |
| 17286 |                |                                                                                                                                                                                                                                                                                        |
| 17287 |                |                                                                                                                                                                                                                                                                                        |

17288 The interaction between *setitimer()* and any of *alarm()*, *sleep()*, or *usleep()* is unspecified.

17289 **RETURN VALUE**

17290 Upon successful completion, *getitimer()* or *setitimer()* shall return 0; otherwise, -1 shall be  
 17291 returned and *errno* set to indicate the error.

17292 **ERRORS**

17293 The *setitimer()* function shall fail if:

|       |          |                                                                                                                                                                                                        |
|-------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 17294 | [EINVAL] | The <i>value</i> argument is not in canonical form. (In canonical form, the number of microseconds is a non-negative integer less than 1,000,000 and the number of seconds is a non-negative integer.) |
| 17295 |          |                                                                                                                                                                                                        |
| 17296 |          |                                                                                                                                                                                                        |

17297 The *getitimer()* and *setitimer()* functions may fail if:

|       |          |                                              |
|-------|----------|----------------------------------------------|
| 17298 | [EINVAL] | The <i>which</i> argument is not recognized. |
|-------|----------|----------------------------------------------|

17299 **EXAMPLES**

17300 None.

17301 **APPLICATION USAGE**

17302 None.

17303 **RATIONALE**

17304 None.

17305 **FUTURE DIRECTIONS**

17306 None.

17307 **SEE ALSO**17308 *alarm()*, *sleep()*, *timer\_getoverrun()*, *ualarm()*, *usleep()*, the Base Definitions volume of

17309 IEEE Std. 1003.1-200x, &lt;signal.h&gt;, &lt;sys/time.h&gt;

17310 **CHANGE HISTORY**

17311 First released in Issue 4, Version 2.

17312 **Issue 5**

17313 Moved from X/OPEN UNIX extension to BASE.

17314 **Issue 6**17315 The **restrict** keyword is added to the *setitimer()* prototype for alignment with the

17316 ISO/IEC 9899:1999 standard.

17317 **NAME**

17318           getlogin, getlogin\_r — get login name

17319 **SYNOPSIS**

17320           #include &lt;unistd.h&gt;

17321           char \*getlogin(void);

17322 TSF       int getlogin\_r(char \*name, size\_t namesize);

17323

17324 **DESCRIPTION**

17325       The *getlogin()* function shall return a pointer to a string containing the user name associated by  
 17326       the login activity with the controlling terminal of the current process. If *getlogin()* returns a non-  
 17327       null pointer, then that pointer points to the name that the user logged in under, even if there are  
 17328       several login names with the same user ID.

17329       The *getlogin()* function need not be reentrant. A function that is not required to be reentrant is  
 17330       not required to be thread-safe.

17331 TSF       The *getlogin\_r()* function puts the name associated by the login activity with the controlling  
 17332       terminal of the current process in the character array pointed to by *name*. The array is *namesize*  
 17333       characters long and should have space for the name and the terminating null character. The  
 17334       maximum size of the login name is {LOGIN\_NAME\_MAX}.

17335       If *getlogin\_r()* is successful, *name* points to the name the user used at login, even if there are  
 17336       several login names with the same user ID.

17337 **RETURN VALUE**

17338       Upon successful completion, *getlogin()* shall return a pointer to the login name or a null pointer  
 17339       if the user's login name cannot be found. Otherwise, it shall return a null pointer and set *errno* to  
 17340       indicate the error.

17341       The return value from *getlogin()* may point to static data whose content is overwritten by each  
 17342       call.

17343 TSF       If successful, the *getlogin\_r()* function shall return zero; otherwise, an error number shall be  
 17344       returned to indicate the error.

17345 **ERRORS**17346       The *getlogin()* and *getlogin\_r()* functions may fail if:

17347       [EMFILE]       {OPEN\_MAX} file descriptors are currently open in the calling process.

17348       [ENFILE]       The maximum allowable number of files is currently open in the system.

17349       [ENXIO]       The calling process has no controlling terminal.

17350       The *getlogin\_r()* function may fail if:

17351 TSF       [ERANGE]       The value of *namesize* is smaller than the length of the string to be returned  
 17352       including the terminating null character.

## 17353 EXAMPLES

17354 **Getting the User Login Name**

17355 The following example calls the *getlogin()* function to obtain the name of the user associated  
 17356 with the calling process, and passes this information to the *getpwnam()* function to get the  
 17357 associated user database information.

```

17358 #include <unistd.h>
17359 #include <sys/types.h>
17360 #include <pwd.h>
17361 #include <stdio.h>
17362 ...
17363 char *lgn;
17364 struct passwd *pw;
17365 ...
17366 if ((lgn = getlogin()) == NULL || (pw = getpwnam(lgn)) == NULL) {
17367 fprintf(stderr, "Get of user information failed.\n"); exit(1);
17368 }
```

17369 **APPLICATION USAGE**

17370 Three names associated with the current process can be determined: *getpwuid(geteuid())* shall  
 17371 return the name associated with the effective user ID of the process; *getlogin()* shall return the  
 17372 name associated with the current login activity; and *getpwuid(getuid())* shall return the name  
 17373 associated with the real user ID of the process.

17374 The *getlogin\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
 17375 of possibly using a static data area that may be overwritten by each call.

17376 **RATIONALE**

17377 The *getlogin()* function returns a pointer to the user's login name. The same user ID may be  
 17378 shared by several login names. If it is desired to get the user database entry that is used during  
 17379 login, the result of *getlogin()* should be used to provide the argument to the *getpwnam()*  
 17380 function. (This might be used to determine the user's login shell, particularly where a single user  
 17381 has multiple login shells with distinct login names, but the same user ID.)

17382 The information provided by the *cuserid()* function, which was originally defined in the  
 17383 POSIX.1-1988 standard and subsequently removed, can be obtained by the following:

```
17384 getpwuid(geteuid())
```

17385 while the information provided by historical implementations of *cuserid()* can be obtained by:

```
17386 getpwuid(getuid())
```

17387 The thread-safe version of this function places the user name in a user-supplied buffer and  
 17388 returns a non-zero value if it fails. The non-thread-safe version may return the name in a static  
 17389 data area that may be overwritten by each call.

17390 **FUTURE DIRECTIONS**

17391 None.

17392 **SEE ALSO**

17393 *getpwnam()*, *getpwuid()*, *geteuid()*, *getuid()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
 17394 <limits.h>, <unistd.h>

17395 **CHANGE HISTORY**

17396 First released in Issue 1. Derived from System V Release 2.0.

17397 **Issue 4**

17398 The `<unistd.h>` header is added to the SYNOPSIS section.

17399 In the RETURN VALUE section, reference to the setting of *errno* is marked as an extension.

17400 The behavior of the function when the login name cannot be found is included in the RETURN  
17401 VALUE section instead of the DESCRIPTION.

17402 The errors [EMFILE], [ENFILE], and [ENXIO] are marked as extensions.

17403 The APPLICATION USAGE section is changed to refer to *getpwuid()* rather than *cuserid()*, which  
17404 may be withdrawn in a future version.

17405 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 17406 • The argument list is explicitly defined as **void**.
- 17407 • The DESCRIPTION is updated to state explicitly that the return value is a pointer to a string  
17408 giving the user name, rather than simply a pointer to the user name as stated in previous  
17409 issues.

17410 **Issue 5**

17411 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
17412 VALUE section.

17413 The *getlogin\_r()* function is included for alignment with the POSIX Threads Extension.

17414 A note indicating that the *getlogin()* function need not be reentrant is added to the  
17415 DESCRIPTION.

17416 **Issue 6**

17417 The *getlogin\_r()* function is marked as part of the Thread-Safe Functions option.

17418 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

17419 The following new requirements on POSIX implementations derive from alignment with the  
17420 Single UNIX Specification:

- 17421 • In the RETURN VALUE section, the requirement to set *errno* on error is added.
- 17422 • The [EMFILE], [ENFILE], and [ENXIO] optional error conditions are added.

17423 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
17424 its avoidance of possibly using a static data area.

## 17425 NAME

17426 getmsg, getpmsg — receive next message from a STREAMS file (STREAMS)

## 17427 SYNOPSIS

17428 XSR #include <stropts.h>

```
17429 int getmsg(int fildes, struct strbuf *restrict ctlptr,
17430 struct strbuf *restrict dataptr, int *restrict flagsp);
17431 int getpmsg(int fildes, struct strbuf *restrict ctlptr,
17432 struct strbuf *restrict dataptr, int *restrict bandp,
17433 int *restrict flagsp);
17434
```

## 17435 DESCRIPTION

17436 The *getmsg()* function shall retrieve the contents of a message located at the head of the  
 17437 STREAM head read queue associated with a STREAMS file and place the contents into one or  
 17438 more buffers. The message contains either a data part, a control part, or both. The data and  
 17439 control parts of the message are placed into separate buffers, as described below. The semantics  
 17440 of each part are defined by the originator of the message.

17441 The *getpmsg()* function does the same thing as *getmsg()*, but provides finer control over the  
 17442 priority of the messages received. Except where noted, all requirements on *getmsg()* also pertain  
 17443 to *getpmsg()*.

17444 The *fildes* argument specifies a file descriptor referencing a STREAMS-based file.

17445 The *ctlptr* and *dataptr* arguments each point to a **strbuf** structure, in which the *buf* member points  
 17446 to a buffer in which the data or control information is to be placed, and the *maxlen* member  
 17447 indicates the maximum number of bytes this buffer can hold. On return, the *len* member  
 17448 contains the number of bytes of data or control information actually received. The *len* member is  
 17449 set to 0 if there is a zero-length control or data part and *len* is set to -1 if no data or control  
 17450 information is present in the message.

17451 When *getmsg()* is called, *flagsp* should point to an integer that indicates the type of message the  
 17452 process is able to receive. This is described further below.

17453 The *ctlptr* argument is used to hold the control part of the message, and *dataptr* is used to hold  
 17454 the data part of the message. If *ctlptr* (or *dataptr*) is a null pointer or the *maxlen* member is -1, the  
 17455 control (or data) part of the message is not processed and is left on the STREAM head read  
 17456 queue, and if the *ctlptr* (or *dataptr*) is not a null pointer, *len* is set to -1. If the *maxlen* member is  
 17457 set to 0 and there is a zero-length control (or data) part, that zero-length part is removed from  
 17458 the read queue and *len* is set to 0. If the *maxlen* member is set to 0 and there are more than 0 bytes  
 17459 of control (or data) information, that information is left on the read queue and *len* is set to 0. If  
 17460 the *maxlen* member in *ctlptr* (or *dataptr*) is less than the control (or data) part of the message,  
 17461 *maxlen* bytes are retrieved. In this case, the remainder of the message is left on the STREAM head  
 17462 read queue and a non-zero return value is provided.

17463 By default, *getmsg()* processes the first available message on the STREAM head read queue.  
 17464 However, a process may choose to retrieve only high-priority messages by setting the integer  
 17465 pointed to by *flagsp* to RS\_HIPRI. In this case, *getmsg()* shall only process the next message if it is  
 17466 a high-priority message. When the integer pointed to by *flagsp* is 0, any message shall be  
 17467 retrieved. In this case, on return, the integer pointed to by *flagsp* is set to RS\_HIPRI if a high-  
 17468 priority message was retrieved, or 0 otherwise.

17469 For *getpmsg()*, the flags are different. The *flagsp* argument points to a bitmask with the following  
 17470 mutually-exclusive flags defined: MSG\_HIPRI, MSG\_BAND, and MSG\_ANY. Like *getmsg()*,  
 17471 *getpmsg()* processes the first available message on the STREAM head read queue. A process may

17472 choose to retrieve only high-priority messages by setting the integer pointed to by *flagsp* to  
 17473 MSG\_HIPRI and the integer pointed to by *bandp* to 0. In this case, *getpmsg()* shall only process  
 17474 the next message if it is a high-priority message. In a similar manner, a process may choose to  
 17475 retrieve a message from a particular priority band by setting the integer pointed to by *flagsp* to  
 17476 MSG\_BAND and the integer pointed to by *bandp* to the priority band of interest. In this case,  
 17477 *getpmsg()* shall only process the next message if it is in a priority band equal to, or greater than,  
 17478 the integer pointed to by *bandp*, or if it is a high-priority message. If a process just wants to get  
 17479 the first message off the queue, the integer pointed to by *flagsp* should be set to MSG\_ANY and  
 17480 the integer pointed to by *bandp* should be set to 0. On return, if the message retrieved was a  
 17481 high-priority message, the integer pointed to by *flagsp* is set to MSG\_HIPRI and the integer  
 17482 pointed to by *bandp* shall be set to 0. Otherwise, the integer pointed to by *flagsp* shall be set to  
 17483 MSG\_BAND and the integer pointed to by *bandp* shall be set to the priority band of the message.

17484 If O\_NONBLOCK is not set, *getmsg()* and *getpmsg()* shall block until a message of the type  
 17485 specified by *flagsp* is available at the front of the STREAM head read queue. If O\_NONBLOCK is  
 17486 set and a message of the specified type is not present at the front of the read queue, *getmsg()* and  
 17487 *getpmsg()* shall fail and set *errno* to [EAGAIN].

17488 If a hangup occurs on the STREAM from which messages are retrieved, *getmsg()* and *getpmsg()*  
 17489 continue to operate normally, as described above, until the STREAM head read queue is empty.  
 17490 Thereafter, they return 0 in the *len* members of *ctlptr* and *dataptr*.

#### 17491 RETURN VALUE

17492 Upon successful completion, *getmsg()* and *getpmsg()* shall return a non-negative value. A value  
 17493 of 0 indicates that a full message was read successfully. A return value of MORECTL indicates  
 17494 that more control information is waiting for retrieval. A return value of MOREDATA indicates  
 17495 that more data is waiting for retrieval. A return value of the bitwise-logical OR of MORECTL  
 17496 and MOREDATA indicates that both types of information remain. Subsequent *getmsg()* and  
 17497 *getpmsg()* calls shall retrieve the remainder of the message. However, if a message of higher  
 17498 priority has come in on the STREAM head read queue, the next call to *getmsg()* or *getpmsg()*  
 17499 shall retrieve that higher-priority message before retrieving the remainder of the previous  
 17500 message.

17501 If the high priority control part of the message is consumed, the message shall be placed back on  
 17502 the queue as a normal message of band 0. Subsequent *getmsg()* and *getpmsg()* calls shall retrieve  
 17503 the remainder of the message. If, however, a priority message arrives or already exists on the  
 17504 STREAM head, the subsequent call to *getmsg()* or *getpmsg()* retrieves the higher-priority  
 17505 message before retrieving the remainder of the message that was put back.

17506 Upon failure, *getmsg()* and *getpmsg()* shall return -1 and set *errno* to indicate the error.

#### 17507 ERRORS

17508 The *getmsg()* and *getpmsg()* functions shall fail if:

|       |           |                                                                                                                                                                                                |  |
|-------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| 17509 | [EAGAIN]  | The O_NONBLOCK flag is set and no messages are available.                                                                                                                                      |  |
| 17510 | [EBADF]   | The <i>fildev</i> argument is not a valid file descriptor open for reading.                                                                                                                    |  |
| 17511 | [EBADMSG] | The queued message to be read is not valid for <i>getmsg()</i> or <i>getpmsg()</i> or a<br>17512 pending file descriptor is at the STREAM head.                                                |  |
| 17513 | [EINTR]   | A signal was caught during <i>getmsg()</i> or <i>getpmsg()</i> .                                                                                                                               |  |
| 17514 | [EINVAL]  | An illegal value was specified by <i>flagsp</i> , or the STREAM or multiplexer<br>17515 referenced by <i>fildev</i> is linked (directly or indirectly) downstream from a<br>17516 multiplexer. |  |

17517 [ENOSTR] A STREAM is not associated with *fildev*.

17518 In addition, *getmsg()* and *getpmsg()* shall fail if the STREAM head had processed an  
17519 asynchronous error before the call. In this case, the value of *errno* does not reflect the result of  
17520 *getmsg()* or *getpmsg()* but reflects the prior error.

## 17521 EXAMPLES

### 17522 Getting Any Message

17523 In the following example, the value of *fd* is assumed to refer to an open STREAMS file. The call  
17524 to *getmsg()* retrieves any available message on the associated STREAM-head read queue,  
17525 returning control and data information to the buffers pointed to by *ctrlbuf* and *databuf*,  
17526 respectively.

```
17527 #include <stropts.h>
17528 ...
17529 int fd;
17530 char ctrlbuf[128];
17531 char databuf[512];
17532 struct strbuf ctrl;
17533 struct strbuf data;
17534 int flags = 0;
17535 int ret;

17536 ctrl.buf = ctrlbuf;
17537 ctrl.maxlen = sizeof(ctrlbuf);

17538 data.buf = databuf;
17539 data.maxlen = sizeof(databuf);

17540 ret = getmsg (fd, &ctrl, &data, &flags);
```

### 17541 Getting the First Message off the Queue

17542 In the following example, the call to *getpmsg()* retrieves the first available message on the  
17543 associated STREAM-head read queue.

```
17544 #include <stropts.h>
17545 ...

17546 int fd;
17547 char ctrlbuf[128];
17548 char databuf[512];
17549 struct strbuf ctrl;
17550 struct strbuf data;
17551 int band = 0;
17552 int flags = MSG_ANY;
17553 int ret;

17554 ctrl.buf = ctrlbuf;
17555 ctrl.maxlen = sizeof(ctrlbuf);

17556 data.buf = databuf;
17557 data.maxlen = sizeof(databuf);

17558 ret = getpmsg (fd, &ctrl, &data, &band, &flags);
```

17559 **APPLICATION USAGE**

17560 None.

17561 **RATIONALE**

17562 None.

17563 **FUTURE DIRECTIONS**

17564 None.

17565 **SEE ALSO**

17566 *poll()*, *putmsg()*, *read()*, *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
17567 `<stropts.h>`, Section 2.6 (on page 539)

17568 **CHANGE HISTORY**

17569 First released in Issue 4, Version 2.

17570 **Issue 5**

17571 Moved from X/OPEN UNIX extension to BASE.

17572 A paragraph regarding “high-priority control parts of messages” is added to the RETURN  
17573 VALUE section.

17574 **Issue 6**

17575 This function is marked as part of the XSI STREAMS Option Group.

17576 The **restrict** keyword is added to the *getmsg()* and *getpmsg()* prototypes for alignment with the  
17577 ISO/IEC 9899:1999 standard.

17578 **NAME**

17579 getnameinfo — get name information

17580 **Notes to Reviewers**17581 *This section with side shading will not appear in the final copy. - Ed.*17582 The IPv6 developers believe that `getnameinfo()` should not truncate the result if the user buffer is  
17583 too small, but return an error status.17584 **SYNOPSIS**

17585 #include &lt;sys/socket.h&gt;

17586 #include &lt;netdb.h&gt;

17587 int getnameinfo(const struct sockaddr \*restrict sa, socklen\_t salen,  
17588 char \*restrict node, socklen\_t nodelen, char \*restrict service,  
17589 socklen\_t servicelen, unsigned flags);17590 **DESCRIPTION**17591 The `getnameinfo()` function translates a socket address to a node name and service location, all of  
17592 which are defined as in `getaddrinfo()`.17593 The `sa` argument points to a socket address structure to be translated.17594 If the `node` argument is non-NULL and the `nodelen` argument is non-zero, then the `node` argument  
17595 points to a buffer able to contain up to `nodelen` characters that receives the node name as a null-  
17596 terminated string. If the `node` argument is NULL or the `nodelen` argument is zero, the node name  
17597 shall not be returned. If the node's name cannot be located, the numeric form of the node's  
17598 address is returned instead of its name.17599 If the `service` argument is non-NULL and the `servicelen` argument is non-zero, then the `service`  
17600 argument points to a buffer able to contain up to `servicelen` characters that receives the service  
17601 name as a null-terminated string. If the `service` argument is NULL or the `servicelen` argument is  
17602 zero, the service name shall not be returned. If the service's name cannot be located, the numeric  
17603 form of the service address (for example, its port number) is returned instead of its name.17604 The `node` and `service` arguments cannot both be NULL.17605 The `flags` argument is a flag that changes the default actions of the function. By default the fully-  
17606 qualified domain name (FQDN) for the host is returned, but:

- 17607 • If the flag bit `NI_NOFQDN` is set, only the node name portion of the FQDN shall be returned  
17608 for local hosts.
- 17609 • If the flag bit `NI_NUMERICHOST` is set, the numeric form of the host's address shall be  
17610 returned instead of its name, under all circumstances.
- 17611 • If the flag bit `NI_NAMEREQD` is set, an error shall be returned if the host's name cannot be  
17612 located.
- 17613 • If the flag bit `NI_NUMERICSERV` is set, the numeric form of the service address shall be  
17614 returned (for example, its port number) instead of its name, under all circumstances.
- 17615 • If the flag bit `NI_DGRAM` is set, this indicates that the service is a datagram service  
17616 (`SOCK_DGRAM`). The default behavior shall assume that the service is a stream service  
17617 (`SOCK_STREAM`).

17618 **Notes:**

- 17619 1. The two `NI_NUMERICxxx` flags are required to support the `-n` flag that many  
17620 commands provide.

17621                   2. The NI\_DGRAM flag is required for the few AF\_INET and AF\_INET6 port  
17622                   numbers (for example, 512-514) that represent different services for UDP and  
17623                   TCP.

17624                The *getnameinfo()* function shall be thread-safe.

#### 17625 RETURN VALUE

17626                A zero return value for *getnameinfo()* indicates successful completion; a non-zero return value  
17627                indicates failure. The possible values for the failures are listed in the ERRORS section.

17628                Upon successful completion, *getnameinfo()* shall return the *node* and *service* names, if requested,  
17629                in the buffers provided. The returned names are always null-terminated strings, and may be  
17630                truncated if the actual values are longer than can be stored in the buffers provided.

#### 17631 ERRORS

17632                The *getnameinfo()* function shall fail and return the corresponding value if:

17633                [EAI\_AGAIN]    The name could not be resolved at this time. Future attempts may succeed.

17634                [EAI\_BADFLAGS]

17635                                The *flags* had an invalid value.

17636                [EAI\_FAIL]     A non-recoverable error occurred.

17637                [EAI\_FAMILY]   The address family was not recognized or the address length was invalid for  
17638                the specified family.

17639                [EAI\_MEMORY]   There was a memory allocation failure.

17640                [EAI\_NONAME]   The name does not resolve for the supplied parameters.

17641                                NI\_NAMEREQD is set and the host's name cannot be located, or both  
17642                                *nodename* and *servname* were null.

17643                [EAI\_SYSTEM]   A system error occurred. The error code can be found in *errno*.

#### 17644 EXAMPLES

17645                None.

#### 17646 APPLICATION USAGE

17647                If the returned values are to be used as part of any further name resolution (for example, passed  
17648                to *getaddrinfo()*), applications shall either provide buffers large enough to store any result  
17649                possible on the system or shall check for truncation and handle that case appropriately.

#### 17650 RATIONALE

17651                None.

#### 17652 FUTURE DIRECTIONS

17653                None.

#### 17654 SEE ALSO

17655                *getaddrinfo()*, *getservbyname()*, *getservbyport()*, *inet\_ntop()*, *socket()*, the Base Definitions volume  
17656                of IEEE Std. 1003.1-200x, <netdb.h>, <sys/socket.h>

#### 17657 CHANGE HISTORY

17658                First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

17659                The **restrict** keyword is added to the *getnameinfo()* prototype for alignment with the  
17660                ISO/IEC 9899:1999 standard.

17661 **NAME**

17662       getnetbyaddr — network database functions

17663 **SYNOPSIS**

17664       #include <netdb.h>

17665       struct netent \*getnetbyaddr(uint32\_t net, int type);

17666 **DESCRIPTION**

17667       Refer to *endnetent()*.

17668 **NAME**

17669       getnetbyname — network database functions

17670 **SYNOPSIS**

17671       #include <netdb.h>

17672       struct netent \*getnetbyname(const char \*name);

17673 **DESCRIPTION**

17674       Refer to *endnetent()*.

17675 **NAME**

17676       getnetent — network database functions

17677 **SYNOPSIS**

17678       #include <netdb.h>

17679       struct netent \*getnetent(void);

17680 **DESCRIPTION**

17681       Refer to *endnetent()*.

17682 **NAME**

17683            getopt, optarg, opterr, optind, optopt — command option parsing

17684 **SYNOPSIS**

17685            #include &lt;unistd.h&gt;

17686            int getopt(int argc, char \* const argv[], const char \*optstring);

17687            extern char \*optarg;

17688            extern int optind, opterr, optopt;

17689 **DESCRIPTION**

17690            The *getopt()* function is a command-line parser that can be used by applications that follow  
 17691            Utility Syntax Guidelines 3, 4, 5, 6, 7, 9, and 10 in the Base Definitions volume of  
 17692            IEEE Std. 1003.1-200x, Section 12.2, Utility Syntax Guidelines. The remaining guidelines are not  
 17693            addressed by *getopt()* and are the responsibility of the application.

17694            The parameters *argc* and *argv* are the argument count and argument array as passed to *main()*  
 17695            (see *exec*). The argument *optstring* is a string of recognized option characters; if a character is  
 17696            followed by a colon, the option takes an argument. All option characters allowed by Utility  
 17697            Syntax Guideline 3 are allowed in *optstring*. The implementation may accept other characters as  
 17698            an extension.

17699            The variable *optind* is the index of the next element of the *argv[]* vector to be processed. It is  
 17700            initialized to 1 by the system, and *getopt()* updates it when it finishes with each element of  
 17701            *argv[]*. When an element of *argv[]* contains multiple option characters, it is unspecified how  
 17702            *getopt()* determines which options have already been processed.

17703            The *getopt()* function shall return the next option character (if one is found) from *argv* that  
 17704            matches a character in *optstring*, if there is one that matches. If the option takes an argument,  
 17705            *getopt()* shall set the variable *optarg* to point to the option-argument as follows:

- 17706            1. If the option was the last character in the string pointed to by an element of *argv*, then  
 17707            *optarg* contains the next element of *argv*, and *optind* is incremented by 2. If the resulting  
 17708            value of *optind* is greater than *argc*, this indicates a missing option-argument, and *getopt()*  
 17709            shall return an error indication.
- 17710            2. Otherwise, *optarg* points to the string following the option character in that element of  
 17711            *argv*, and *optind* is incremented by 1.

17712            If, when *getopt()* is called:

17713            *argv[optind]*    is a null pointer  
 17714            \**argv[optind]*    is not the character *-*  
 17715            *argv[optind]*    points to the string *"-"*

17716            *getopt()* shall return *-1* without changing *optind*. If:

17717            *argv[optind]*    points to the string *"--"*

17718            *getopt()* shall return *-1* after incrementing *optind*.

17719            If *getopt()* encounters an option character that is not contained in *optstring*, it shall return the  
 17720            question-mark ('?') character. If it detects a missing option-argument, it shall return the colon  
 17721            character (':') if the first character of *optstring* was a colon, or a question-mark character ('?')  
 17722            otherwise. In either case, *getopt()* shall set the variable *optopt* to the option character that caused  
 17723            the error. If the application has not set the variable *opterr* to 0 and the first character of *optstring*  
 17724            is not a colon, *getopt()* shall also print a diagnostic message to *stderr* in the format specified for  
 17725            the *getopts* utility.

17726 The *getopt()* function need not be reentrant. A function that is not required to be reentrant is not  
17727 required to be thread-safe.

#### 17728 RETURN VALUE

17729 The *getopt()* function shall return the next option character specified on the command line.

17730 A colon (':') shall be returned if *getopt()* detects a missing argument and the first character of  
17731 *optstring* was a colon (':').

17732 A question mark ('?') shall be returned if *getopt()* encounters an option character not in  
17733 *optstring* or detects a missing argument and the first character of *optstring* was not a colon (':').

17734 Otherwise, *getopt()* shall return -1 when all command line options are parsed.

#### 17735 ERRORS

17736 No errors are defined.

#### 17737 EXAMPLES

##### 17738 Parsing Command Line Options

17739 The following code fragment shows how you might process the arguments for a utility that can  
17740 take the mutually-exclusive options *a* and *b* and the options *f* and *o*, both of which require  
17741 arguments:

```
17742 #include <unistd.h>
17743 int
17744 main(int argc, char *argv[])
17745 {
17746 int c;
17747 int bflg, aflag, errflag;
17748 char *ifile;
17749 char *ofile;
17750 extern char *optarg;
17751 extern int optind, optopt;
17752 . . .
17753 while ((c = getopt(argc, argv, ":abf:o:")) != -1) {
17754 switch(c) {
17755 case 'a':
17756 if (bflg)
17757 errflag++;
17758 else
17759 aflag++;
17760 break;
17761 case 'b':
17762 if (aflag)
17763 errflag++;
17764 else {
17765 bflg++;
17766 bproc();
17767 }
17768 break;
17769 case 'f':
17770 ifile = optarg;
17771 break;
17772 case 'o':
```

```

17773 ofile = optarg;
17774 break;
17775 case ':': /* -f or -o without operand */
17776 fprintf(stderr,
17777 "Option -%c requires an operand\n", optopt);
17778 errflg++;
17779 break;
17780 case '?':
17781 fprintf(stderr,
17782 "Unrecognized option: -%c\n", optopt);
17783 errflg++;
17784 }
17785 }
17786 if (errflg) {
17787 fprintf(stderr, "usage: . . . ");
17788 exit(2);
17789 }
17790 for (; optind < argc; optind++) {
17791 if (access(argv[optind], R_OK)) {
17792 . . .
17793 }

```

17794 This code accepts any of the following as equivalent:

```

17795 cmd -ao arg path path
17796 cmd -a -o arg path path
17797 cmd -o arg -a path path
17798 cmd -a -o arg — path path
17799 cmd -a -oarg path path
17800 cmd -aoarg path path

```

### 17801 **Checking Options and Arguments**

17802 The following example parses a set of command line options and prints messages to standard  
17803 output for each option and argument that it encounters.

```

17804 #include <unistd.h>
17805 #include <stdio.h>
17806 ...
17807 int c;
17808 char *filename;
17809 extern char *optarg;
17810 extern int optind, optopt, opterr;
17811 ...
17812 while ((c = getopt(argc, argv, ":abf:")) != -1) {
17813 switch(c) {
17814 case 'a':
17815 printf("a is set\n");
17816 break;
17817 case 'b':
17818 printf("b is set\n");
17819 break;
17820 case 'f':
17821 filename = optarg;

```

```

17822 printf("filename is %s\n", filename);
17823 break;
17824 case ':':
17825 printf("--%c without filename\n", optopt);
17826 break;
17827 case '?':
17828 printf("unknown arg %c\n", optopt);
17829 break;
17830 }
17831 }

```

### 17832 **Selecting Options from the Command Line**

17833 The following example selects the type of database routines the user wants to use based on the  
 17834 *Options* argument.

```

17835 #include <unistd.h>
17836 #include <string.h>
17837 ...
17838 char *Options = "hdbtl";
17839 ...
17840 int dbtype, i;
17841 char c;
17842 char *st;
17843 ...
17844 dbtype = 0;
17845 while ((c = getopt(argc, argv, Options)) != -1) {
17846 if ((st = strchr(Options, c)) != NULL) {
17847 dbtype = st - Options;
17848 break;
17849 }
17850 }

```

### 17851 **APPLICATION USAGE**

17852 The *getopt()* function is only required to support option characters included in Utility Syntax  
 17853 Guideline 3. Many historical implementations of *getopt()* support other characters as options.  
 17854 This is an allowed extension, but applications that use extensions are not maximally portable.  
 17855 Note that support for multi-byte option characters is only possible when such characters can be  
 17856 represented as type **int**.

### 17857 **RATIONALE**

17858 The *optopt* variable represents historical practice and allows the application to obtain the identity  
 17859 of the invalid option.

17860 The description has been written to make it clear that *getopt()*, like the *getopts* utility, deals with  
 17861 option-arguments whether separated from the option by <blank> characters or not. Note that  
 17862 the requirements on *getopt()* and *getopts* are more stringent than the Utility Syntax Guidelines.

17863 The *getopt()* function shall return  $-1$ , rather than EOF, so that <**stdio.h**> is not required.

17864 The special significance of a colon as the first character of *optstring* makes *getopt()* consistent  
 17865 with the *getopts* utility. It allows an application to make a distinction between a missing  
 17866 argument and an incorrect option letter without having to examine the option letter. It is true  
 17867 that a missing argument can only be detected in one case, but that is a case that has to be  
 17868 considered.

17869 **FUTURE DIRECTIONS**

17870 None.

17871 **SEE ALSO**

17872 *exec*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>, the Shell and Utilities  
17873 volume of IEEE Std. 1003.1-200x

17874 **CHANGE HISTORY**

17875 First released in Issue 1. Derived from Issue 1 of the SVID.

17876 **Issue 4**

17877 The following changes are incorporated for alignment with the ISO POSIX-2 standard:

- 17878 • The <**unistd.h**> header is added to the SYNOPSIS section and <**stdio.h**> is deleted.
- 17879 • The type of argument *argv* is changed from **char\*\*** to **char\*const[ ]**.
- 17880 • The integer *optopt* is added to the list of external data items.
- 17881 • The DESCRIPTION is largely rewritten, without functional change, for alignment with the  
17882 ISO POSIX-2 standard, although the following differences should be noted:
  - 17883 — If the function detects a missing option-argument, it returns a colon ( ':' ) and sets *optopt*  
17884 to the option character.
  - 17885 — The termination conditions under which *getopt()* returns -1 are extended. Also note that  
17886 the termination condition is explicitly -1, rather than the value of EOF.
- 17887 • The EXAMPLES section is changed to illustrate the new functionality.

17888 **Issue 5**17889 A note indicating that the *getopt()* function need not be reentrant is added to the DESCRIPTION.17890 **Issue 6**

17891 IEEE PASC Interpretation 1003.2 #150 is applied.

17892 **NAME**

17893 getpeername — get the name of the peer socket

17894 **SYNOPSIS**

17895 #include <sys/socket.h>

17896 int getpeername(int *socket*, struct sockaddr \*restrict *address*,  
17897 socklen\_t \**address\_len*);

17898 **DESCRIPTION**

17899 The *getpeername()* function shall retrieve the peer address of the specified socket, store this  
17900 address in the **sockaddr** structure pointed to by the *address* argument, and store the length of this  
17901 address in the object pointed to by the *address\_len* argument.

17902 If the actual length of the address is greater than the length of the supplied **sockaddr** structure,  
17903 the stored address shall be truncated.

17904 If the protocol permits connections by unbound clients, and the peer is not bound, then the value  
17905 stored in the object pointed to by *address* is unspecified.

17906 **RETURN VALUE**

17907 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
17908 indicate the error.

17909 **ERRORS**

17910 The *getpeername()* function shall fail if:

17911 [EBADF] The *socket* argument is not a valid file descriptor.

17912 [EINVAL] The socket has been shut down.

17913 [ENOTCONN] The socket is not connected or otherwise has not had the peer prespecified.

17914 [ENOTSOCK] The *socket* argument does not refer to a socket.

17915 [EOPNOTSUPP] The operation is not supported for the socket protocol.

17916 The *getpeername()* function may fail if:

17917 [ENOBUFS] Insufficient resources were available in the system to complete the call.

17918 **EXAMPLES**

17919 None.

17920 **APPLICATION USAGE**

17921 None.

17922 **RATIONALE**

17923 None.

17924 **FUTURE DIRECTIONS**

17925 None.

17926 **SEE ALSO**

17927 *accept()*, *bind()*, *getsockname()*, *socket()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
17928 <sys/socket.h>

17929 **CHANGE HISTORY**

17930 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

17931 The **restrict** keyword is added to the *getpeername()* prototype for alignment with the  
17932 ISO/IEC 9899:1999 standard.

17933 **NAME**

17934 getpgid — get the process group ID for a process

17935 **SYNOPSIS**

17936 xSI #include <unistd.h>

17937 pid\_t getpgid(pid\_t pid);

17938

17939 **DESCRIPTION**

17940 The *getpgid()* function shall return the process group ID of the process whose process ID is equal  
17941 to *pid*. If *pid* is equal to 0, *getpgid()* shall return the process group ID of the calling process.

17942 **RETURN VALUE**

17943 Upon successful completion, *getpgid()* shall return a process group ID. Otherwise, it shall return  
17944 (**pid\_t**)-1 and set *errno* to indicate the error.

17945 **ERRORS**

17946 The *getpgid()* function shall fail if:

17947 [EPERM] The process whose process ID is equal to *pid* is not in the same session as the  
17948 calling process, and the implementation does not allow access to the process  
17949 group ID of that process from the calling process.

17950 [ESRCH] There is no process with a process ID equal to *pid*.

17951 The *getpgid()* function may fail if:

17952 [EINVAL] The value of the *pid* argument is invalid.

17953 **EXAMPLES**

17954 None.

17955 **APPLICATION USAGE**

17956 None.

17957 **RATIONALE**

17958 None.

17959 **FUTURE DIRECTIONS**

17960 None.

17961 **SEE ALSO**

17962 *exec*, *fork()*, *getpgrp()*, *getpid()*, *getsid()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
17963 IEEE Std. 1003.1-200x, <unistd.h>

17964 **CHANGE HISTORY**

17965 First released in Issue 4, Version 2.

17966 **Issue 5**

17967 Moved from X/OPEN UNIX extension to BASE.

17968 **NAME**

17969           getpgrp — get the process group ID of the calling process

17970 **SYNOPSIS**

17971           #include <unistd.h>

17972           pid\_t getpgrp(void);

17973 **DESCRIPTION**

17974           The *getpgrp()* function shall return the process group ID of the calling process.

17975 **RETURN VALUE**

17976           The *getpgrp()* function is always successful and no return value is reserved to indicate an error.

17977 **ERRORS**

17978           No errors are defined.

17979 **EXAMPLES**

17980           None.

17981 **APPLICATION USAGE**

17982           None.

17983 **RATIONALE**

17984           4.3 BSD provides a *getpgrp()* function that returns the process group ID for a specified process.  
17985           Although this function is used to support job control, all known job control shells always specify  
17986           the calling process with this function. Thus, the simpler System V *getpgrp()* suffices, and the  
17987           added complexity of the 4.3 BSD *getpgrp()* has been omitted from this volume of  
17988           IEEE Std. 1003.1-200x.

17989 **FUTURE DIRECTIONS**

17990           None.

17991 **SEE ALSO**

17992           *exec*, *fork()*, *getpgid()*, *getpid()*, *getppid()*, *kill()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
17993           IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>

17994 **CHANGE HISTORY**

17995           First released in Issue 1. Derived from Issue 1 of the SVID.

17996 **Issue 4**

17997           The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
17998           XSI-conformant systems.

17999           The <unistd.h> header is added to the SYNOPSIS section.

18000           The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 18001           • The argument list is explicitly defined as **void**.

18002 **Issue 6**

18003           In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

18004           The following new requirements on POSIX implementations derive from alignment with the  
18005           Single UNIX Specification:

- 18006           • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
18007           required for conforming implementations of previous POSIX specifications, it was not  
18008           required for UNIX applications.

18009 **NAME**

18010            *getpid* — get the process ID

18011 **SYNOPSIS**

18012            #include <unistd.h>

18013            pid\_t *getpid*(void);

18014 **DESCRIPTION**

18015            The *getpid*() function shall return the process ID of the calling process.

18016 **RETURN VALUE**

18017            The *getpid*() function is always successful and no return value is reserved to indicate an error.

18018 **ERRORS**

18019            No errors are defined.

18020 **EXAMPLES**

18021            None.

18022 **APPLICATION USAGE**

18023            None.

18024 **RATIONALE**

18025            None.

18026 **FUTURE DIRECTIONS**

18027            None.

18028 **SEE ALSO**

18029            *exec*, *fork*(), *getpgrp*(), *getppid*(), *kill*(), *setpgid*(), *setsid*(), the Base Definitions volume of  
18030            IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>

18031 **CHANGE HISTORY**

18032            First released in Issue 1. Derived from Issue 1 of the SVID.

18033 **Issue 4**

18034            The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
18035            XSI-conformant systems.

18036            The <unistd.h> header is added to the SYNOPSIS section.

18037            The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 18038            • The argument list is explicitly defined as **void**.

18039 **Issue 6**

18040            In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

18041            The following new requirements on POSIX implementations derive from alignment with the  
18042            Single UNIX Specification:

- 18043            • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
18044            required for conforming implementations of previous POSIX specifications, it was not  
18045            required for UNIX applications.

18046 **NAME**

18047       getpmsg — receive next message from a STREAMS file

18048 **SYNOPSIS**

18049 XSI     #include <stropts.h>

```
18050 int getpmsg(int fildes, struct strbuf *restrict ctlptr,
18051 struct strbuf *restrict dataptr, int *restrict bandp,
18052 int *restrict flagsp);
```

18053

18054 **DESCRIPTION**

18055       Refer to *getmsg()*.

18056 **NAME**

18057           getppid — get the parent process ID

18058 **SYNOPSIS**

18059           #include <unistd.h>

18060           pid\_t getppid(void);

18061 **DESCRIPTION**

18062           The *getppid()* function shall return the parent process ID of the calling process.

18063 **RETURN VALUE**

18064           The *getppid()* function is always successful and no return value is reserved to indicate an error.

18065 **ERRORS**

18066           No errors are defined.

18067 **EXAMPLES**

18068           None.

18069 **APPLICATION USAGE**

18070           None.

18071 **RATIONALE**

18072           None.

18073 **FUTURE DIRECTIONS**

18074           None.

18075 **SEE ALSO**

18076           *exec*, *fork()*, *getpgid()*, *getpgrp()*, *getpid()*, *kill()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
18077           IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>

18078 **CHANGE HISTORY**

18079           First released in Issue 1. Derived from Issue 1 of the SVID.

18080 **Issue 4**

18081           The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
18082           XSI-conformant systems.

18083           The <unistd.h> header is added to the SYNOPSIS section.

18084           The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 18085           • The argument list is explicitly defined as **void**.

18086 **Issue 6**

18087           In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

18088           The following new requirements on POSIX implementations derive from alignment with the  
18089           Single UNIX Specification:

- 18090           • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
18091           required for conforming implementations of previous POSIX specifications, it was not  
18092           required for UNIX applications.

## 18093 NAME

18094 getpriority, setpriority — get or set the nice value

## 18095 SYNOPSIS

18096 XSI `#include <sys/resource.h>`18097 `int getpriority(int which, id_t who);`18098 `int setpriority(int which, id_t who, int value);`

18099

## 18100 DESCRIPTION

18101 The *getpriority()* function obtains the nice value of a process, process group, or user. The  
 18102 *setpriority()* function sets the nice value of a process, process group, or user to *value*+ {NZERO}.

18103 Target processes are specified by the values of the *which* and *who* arguments. The *which*  
 18104 argument may be one of the following values: PRIO\_PROCESS, PRIO\_PGRP, or PRIO\_USER,  
 18105 indicating that the *who* argument is to be interpreted as a process ID, a process group ID, or an  
 18106 effective user ID, respectively. A 0 value for the *who* argument specifies the current process,  
 18107 process group, or user.

18108 The nice value set with *setpriority()* shall be applied to the process. If the process is multi-  
 18109 threaded, the nice value shall affect all system scope threads in the process.

18110 If more than one process is specified, *getpriority()* shall return value {NZERO} less than the  
 18111 lowest nice value pertaining to any of the specified processes, and *setpriority()* sets the nice  
 18112 values of all of the specified processes to *value*+ {NZERO}.

18113 The default nice value is {NZERO}; lower nice values cause more favorable scheduling. While  
 18114 the range of valid nice values is [0,{NZERO}\*2 -1], implementations may enforce more  
 18115 restrictive limits. If *value*+ {NZERO} is less than the system's lowest supported nice value,  
 18116 *setpriority()* sets the nice value to the lowest supported value; if *value*+ {NZERO} is greater than  
 18117 the system's highest supported nice value, *setpriority()* sets the nice value to the highest  
 18118 supported value.

18119 Only a process with appropriate privileges can lower its nice value.

18120 PS|TPS Any processes or threads using SCHED\_FIFO or SCHED\_RR are unaffected by a call to  
 18121 *setpriority()*. This is not considered an error.

18122 The effect of changing the nice value may vary depending on the process-scheduling algorithm  
 18123 in effect.

18124 Because *getpriority()* can return the value -1 on successful completion, it is necessary to set *errno*  
 18125 to 0 prior to a call to *getpriority()*. If *getpriority()* returns the value -1, then *errno* can be checked  
 18126 to see if an error occurred or if the value is a legitimate nice value.

## 18127 RETURN VALUE

18128 Upon successful completion, *getpriority()* shall return an integer in the range from -{NZERO} to  
 18129 {NZERO}-1. Otherwise, -1 shall be returned and *errno* set to indicate the error.

18130 Upon successful completion, *setpriority()* shall return 0; otherwise, -1 shall be returned and *errno*  
 18131 set to indicate the error.

## 18132 ERRORS

18133 The *getpriority()* and *setpriority()* functions shall fail if:

18134 [ESRCH] No process could be located using the *which* and *who* argument values  
 18135 specified.

18136 [EINVAL] The value of the *which* argument was not recognized, or the value of the *who* |  
 18137 argument is not a valid process ID, process group ID, or user ID.

18138 In addition, *setpriority()* may fail if:

18139 [EPERM] A process was located, but neither the real nor effective user ID of the |  
 18140 executing process match the effective user ID of the process whose nice value |  
 18141 is being changed.

18142 [EACCES] A request was made to change the nice value to a lower numeric value and |  
 18143 the current process does not have appropriate privileges.

#### 18144 EXAMPLES

##### 18145 Using *getpriority()*

18146 The following example returns the current scheduling priority for the process ID returned by the |  
 18147 call to *getpid()*.

```
18148 #include <sys/resource.h>
18149 ...
18150 int which = PRIO_PROCESS;
18151 id_t pid;
18152 int ret;

18153 pid = getpid();
18154 ret = getpriority(which, pid);
```

##### 18155 Using *setpriority()*

18156 The following example sets the priority for the current process ID to  $-20$ .

```
18157 #include <sys/resource.h>
18158 ...
18159 int which = PRIO_PROCESS;
18160 id_t pid;
18161 int priority = -20;
18162 int ret;

18163 pid = getpid();
18164 ret = setpriority(which, pid, priority);
```

#### 18165 APPLICATION USAGE

18166 The *getpriority()* and *setpriority()* functions work with an offset nice value (nice value |  
 18167  $-\{\text{NZERO}\}$ ). The nice value is in the range  $[0, 2*\{\text{NZERO}\} - 1]$ , while the return value for |  
 18168 *getpriority()* and the third parameter for *setpriority()* are in the range  $[-\{\text{NZERO}\}, \{\text{NZERO}\} - 1]$ .

#### 18169 RATIONALE

18170 None.

#### 18171 FUTURE DIRECTIONS

18172 None.

#### 18173 SEE ALSO

18174 *nice()*, *sched\_get\_priority\_max()*, *sched\_setscheduler()*, the Base Definitions volume of |  
 18175 IEEE Std. 1003.1-200x, <sys/resource.h>

18176 **CHANGE HISTORY**

18177 First released in Issue 4, Version 2.

18178 **Issue 5**

18179 Moved from X/OPEN UNIX extension to BASE.

18180 The DESCRIPTION is reworded in terms of the nice value rather than *priority* to avoid confusion  
18181 with functionality in the POSIX Realtime Extension.

18182 **NAME**

18183       getprotobyname — network protocol database functions

18184 **SYNOPSIS**

18185       #include <netdb.h>

18186       struct protoent \*getprotobyname(const char \*name);

18187 **DESCRIPTION**

18188       Refer to *endprotoent()*.

18189 **NAME**

18190           getprotobynumber — network protocol database functions

18191 **SYNOPSIS**

18192           #include <netdb.h>

18193           struct protoent \*getprotobynumber(int *proto*);

18194 **DESCRIPTION**

18195           Refer to *endprotoent()*.

18196 **NAME**

18197       getprotoent — network protocol database functions

18198 **SYNOPSIS**

18199       #include <netdb.h>

18200       struct protoent \*getprotoent(void);

18201 **DESCRIPTION**

18202       Refer to *endprotoent()*.

18203 **NAME**

18204           getpwent — get user database entry

18205 **SYNOPSIS**

18206 xSI       #include <pwd.h>

18207           struct passwd \*getpwent(void);

18208

18209 **DESCRIPTION**

18210           Refer to *endpwent()*.

## 18211 NAME

18212 getpwnam, getpwnam\_r — search user database for a name

## 18213 SYNOPSIS

18214 #include <pwd.h>

18215 struct passwd \*getpwnam(const char \*name);

18216 TSF int getpwnam\_r(const char \*name, struct passwd \*pwd, char \*buffer,

18217 size\_t bufsize, struct passwd \*\*result);

18218

## 18219 DESCRIPTION

18220 The *getpwnam()* function shall search the user database for an entry with a matching *name*.

18221 The *getpwnam()* function need not be reentrant. A function that is not required to be reentrant is  
18222 not required to be thread-safe.

18223 Applications wishing to check for error situations should set *errno* to 0 before calling  
18224 *getpwnam()*. If *getpwnam()* returns a null pointer and *errno* is non-zero, an error occurred.

18225 TSF The *getpwnam\_r()* function updates the **passwd** structure pointed to by *pwd* and stores a pointer  
18226 to that structure at the location pointed to by *result*. The structure shall contain an entry from  
18227 the user database with a matching *name*. Storage referenced by the structure is allocated from  
18228 the memory provided with the *buffer* parameter, which is *bufsize* characters in size.

## 18229 Notes to Reviewers

18230 *This section with side shading will not appear in the final copy. - Ed.*

18231 D3, XSH, ERN 301 says that the size above is in bytes, not characters, and proposes changing  
18232 "characters" to "bytes".

18233 The maximum size needed for this buffer can be determined with the  
18234 `{_SC_GETPW_R_SIZE_MAX} sysconf()` parameter. A NULL pointer shall be returned at the  
18235 location pointed to by *result* on error or if the requested entry is not found.

## 18236 RETURN VALUE

18237 The *getpwnam()* function shall return a pointer to a **struct passwd** with the structure as defined  
18238 in <pwd.h> with a matching entry if found. A null pointer shall be returned if the requested  
18239 entry is not found, or an error occurs. On error, *errno* shall be set to indicate the error.

18240 The return value may point to a static area which is overwritten by a subsequent call to  
18241 *getpwent()*, *getpwnam()*, or *getpwuid()*.

18242 TSF If successful, the *getpwnam\_r()* function shall return zero; otherwise, an error number shall be  
18243 returned to indicate the error.

## 18244 ERRORS

18245 The *getpwnam()* and *getpwnam\_r()* functions may fail if:

18246 [EIO] An I/O error has occurred.

18247 [EINTR] A signal was caught during *getpwnam()*.

18248 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

18249 [ENFILE] The maximum allowable number of files is currently open in the system.

18250 TSF The *getpwnam\_r()* function may fail if:

18251 TSF [ERANGE] Insufficient storage was supplied via *buffer* and *bufsize* to contain the data to  
18252 be referenced by the resulting **passwd** structure.

18253 **EXAMPLES**18254 **Getting an Entry for the Login Name**

18255 The following example uses the *getlogin()* function to return the name of the user who logged in;  
18256 this information is passed to the *getpwnam()* function to get the user database entry for that user.

```
18257 #include <sys/types.h>
18258 #include <pwd.h>
18259 #include <unistd.h>
18260 #include <stdio.h>
18261 #include <stdlib.h>
18262 ...
18263 char *lgn;
18264 struct passwd *pw;
18265 ...
18266 if ((lgn = getlogin()) == NULL || (pw = getpwnam(lgn)) == NULL) {
18267 fprintf(stderr, "Get of user information failed.\n"); exit(1);
18268 }
18269 ...
```

18270 **APPLICATION USAGE**

18271 Three names associated with the current process can be determined: *getpwuid(geteuid())* returns  
18272 the name associated with the effective user ID of the process; *getlogin()* returns the name  
18273 associated with the current login activity; and *getpwuid(getuid())* returns the name associated  
18274 with the real user ID of the process.

18275 The *getpwnam\_r()* function is thread-safe and shall return values in a user-supplied buffer  
18276 instead of possibly using a static data area that may be overwritten by each call.

18277 **RATIONALE**

18278 None.

18279 **FUTURE DIRECTIONS**

18280 None.

18281 **SEE ALSO**

18282 *getpwuid()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<limits.h>**, **<pwd.h>**,  
18283 **<sys/types.h>**

18284 **CHANGE HISTORY**

18285 First released in Issue 1. Derived from System V Release 2.0.

18286 **Issue 4**

18287 The DESCRIPTION is clarified.

18288 The **<sys/types.h>** header is now marked as optional (OH); this header need not be included on  
18289 XSI-conformant systems.

18290 The last sentence in the RETURN VALUE section, indicating that *errno* is set on error, is marked  
18291 as an extension.

18292 The errors [EIO], [EINTR], [EMFILE], and [ENFILE] are marked as extensions.

18293 The APPLICATION USAGE section is expanded to warn about possible reuses of the area used  
18294 to pass the return value, and to indicate how applications should check for errors.

18295 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 18296
- The type of argument *name* is changed from **char\*** to **const char\***.
- 18297 **Issue 5**
- 18298 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
18299 VALUE section.
- 18300 The *getpwnam\_r()* function is included for alignment with the POSIX Threads Extension.
- 18301 A note indicating that the *getpwnam()* function need not be reentrant is added to the  
18302 DESCRIPTION.
- 18303 **Issue 6**
- 18304 The *getpwnam\_r()* function is marked as part of the Thread-Safe Functions option.
- 18305 The Open Group corrigenda item U028/3 has been applied correcting text in the DESCRIPTION  
18306 describing matching the *name*.
- 18307 In the SYNOPSIS, the inclusion of **<sys/types.h>** is no longer required.
- 18308 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.
- 18309 The following new requirements on POSIX implementations derive from alignment with the  
18310 Single UNIX Specification:
- The requirement to include **<sys/types.h>** has been removed. Although **<sys/types.h>** was required for conforming implementations of previous POSIX specifications, it was not required for UNIX applications.
  - In the RETURN VALUE section, the requirement to set *errno* on error is added.
  - The [EMFILE], [ENFILE], and [ENXIO] optional error conditions are added.
- 18316 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
18317 its avoidance of possibly using a static data area.

## 18318 NAME

18319 getpwuid, getpwuid\_r — search user database for a user ID

## 18320 SYNOPSIS

18321 #include <pwd.h>

18322 struct passwd \*getpwuid(uid\_t uid);

18323 TSF int getpwuid\_r(uid\_t uid, struct passwd \*pwd, char \*buffer,

18324 size\_t bufsize, struct passwd \*\*result);

18325

## 18326 DESCRIPTION

18327 The *getpwuid()* function shall search the user database for an entry with a matching *uid*.

18328 The *getpwuid()* function need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

18330 Applications wishing to check for error situations should set *errno* to 0 before calling *getpwuid()*.  
18331 If *getpwuid()* returns a null pointer and *errno* is set to non-zero, an error occurred.

18332 TSF The *getpwuid\_r()* function updates the **passwd** structure pointed to by *pwd* and stores a pointer  
18333 to that structure at the location pointed to by *result*. The structure shall contain an entry from  
18334 the user database with a matching *uid*. Storage referenced by the structure is allocated from the  
18335 memory provided with the *buffer* parameter, which is *bufsize* characters in size. The maximum  
18336 size needed for this buffer can be determined with the `{_SC_GETPW_R_SIZE_MAX}` *sysconf()*  
18337 parameter. A NULL pointer shall be returned at the location pointed to by *result* on error or if the  
18338 requested entry is not found.

## 18339 RETURN VALUE

18340 The *getpwuid()* function shall return a pointer to a **struct passwd** with the structure as defined in  
18341 <**pwd.h**> with a matching entry if found. A null pointer shall be returned if the requested entry  
18342 is not found, or an error occurs. On error, *errno* shall be set to indicate the error.

18343 The return value may point to a static area which is overwritten by a subsequent call to  
18344 *getpwent()*, *getpwnam()*, or *getpwuid()*.

18345 TSF If successful, the *getpwuid\_r()* function shall return zero; otherwise, an error number shall be  
18346 returned to indicate the error.

## 18347 ERRORS

18348 The *getpwuid()* and *getpwuid\_r()* functions may fail if:

18349 [EIO] An I/O error has occurred.

18350 [EINTR] A signal was caught during *getpwuid()*.

18351 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

18352 [ENFILE] The maximum allowable number of files is currently open in the system.

18353 TSF The *getpwuid\_r()* function may fail if:

18354 TSF [ERANGE] Insufficient storage was supplied via *buffer* and *bufsize* to contain the data to  
18355 be referenced by the resulting **passwd** structure.

18356 **EXAMPLES**18357 **Getting an Entry for the Root User**

18358 The following example gets the user database entry for the user with user ID 0 (root).

```
18359 #include <sys/types.h>
18360 #include <pwd.h>
18361 ...
18362 uid_t id = 0;
18363 struct passwd *pwd;
18364 pwd = getpwuid(id);
```

18365 **Finding the Name for the Effective User ID**

18366 The following example defines *pws* as a pointer to a structure of type **passwd**, which is used to  
 18367 store the structure pointer returned by the call to the *getpwuid()* function. The *geteuid()* function  
 18368 shall return the effective user ID of the calling process; this is used as the search criteria for the  
 18369 *getpwuid()* function. The call to *getpwuid()* shall return a pointer to the structure containing that  
 18370 user ID value.

```
18371 #include <unistd.h>
18372 #include <sys/types.h>
18373 #include <pwd.h>
18374 ...
18375 struct passwd *pws;
18376 pws = getpwuid(geteuid());
```

18377 **Finding an Entry in the User Database**

18378 The following example uses *getpwuid()* to search the user database for a user ID that was  
 18379 previously stored in a **stat** structure, then prints out the user name if it is found. If the user is not  
 18380 found, the program prints the numeric value of the user ID for the entry.

```
18381 #include <sys/types.h>
18382 #include <pwd.h>
18383 #include <stdio.h>
18384 ...
18385 struct stat statbuf;
18386 struct passwd *pwd;
18387 ...
18388 if ((pwd = getpwuid(statbuf.st_uid)) != NULL)
18389 printf(" %-8.8s", pwd->pw_name);
18390 else
18391 printf(" %-8d", statbuf.st_uid);
```

18392 **APPLICATION USAGE**

18393 Three names associated with the current process can be determined: *getpwuid(geteuid())* returns  
 18394 the name associated with the effective user ID of the process; *getlogin()* returns the name  
 18395 associated with the current login activity; and *getpwuid(getuid())* returns the name associated  
 18396 with the real user ID of the process.

18397 The *getpwuid\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
 18398 of possibly using a static data area that may be overwritten by each call.

18399 **RATIONALE**

18400 None.

18401 **FUTURE DIRECTIONS**

18402 None.

18403 **SEE ALSO**18404 *getpwnam()*, *geteuid()*, *getuid()*, *getlogin()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
18405 `<limits.h>`, `<pwd.h>`, `<sys/types.h>`18406 **CHANGE HISTORY**

18407 First released in Issue 1. Derived from System V Release 2.0.

18408 **Issue 4**

18409 The DESCRIPTION is clarified.

18410 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on  
18411 XSI-conformant systems.18412 The last sentence in the RETURN VALUE section, indicating that *errno* is set on error, is marked  
18413 as an extension.

18414 The errors [EIO], [EINTR], [EMFILE], and [ENFILE] are marked as extensions.

18415 A note is added to the APPLICATION USAGE section indicating how an application should  
18416 check for errors.18417 **Issue 5**18418 Normative text previously in the APPLICATION USAGE section is moved to the RETURN  
18419 VALUE section.18420 The *getpwuid\_r()* function is included for alignment with the POSIX Threads Extension.18421 A note indicating that the *getpwuid()* function need not be reentrant is added to the  
18422 DESCRIPTION.18423 **Issue 6**18424 The *getpwuid\_r()* function is marked as part of the Thread-Safe Functions option.18425 The Open Group corrigenda item U028/3 has been applied correcting text in the DESCRIPTION  
18426 describing matching the *uid*.18427 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

18428 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

18429 The following new requirements on POSIX implementations derive from alignment with the  
18430 Single UNIX Specification:18431 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
18432 required for conforming implementations of previous POSIX specifications, it was not  
18433 required for UNIX applications.18434 • In the RETURN VALUE section, the requirement to set *errno* on error is added.

18435 • The [EIO], [EINTR], [EMFILE], and [ENFILE] optional error conditions are added.

18436 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
18437 its avoidance of possibly using a static data area.

18438 **NAME**

18439 getrlimit, setrlimit — control maximum resource consumption

18440 **SYNOPSIS**18441 xSI 

```
#include <sys/resource.h>
```

18442 

```
int getrlimit(int resource, struct rlimit *rlp);
```

18443 

```
int setrlimit(int resource, const struct rlimit *rlp);
```

18444

18445 **DESCRIPTION**18446 Limits on the consumption of a variety of resources by the calling process may be obtained with  
18447 *getrlimit()* and set with *setrlimit()*.18448 Each call to either *getrlimit()* or *setrlimit()* identifies a specific resource to be operated upon as  
18449 well as a resource limit. A resource limit is represented by an **rlimit** structure. The *rlim\_cur*  
18450 member specifies the current or soft limit and the *rlim\_max* member specifies the maximum or  
18451 hard limit. Soft limits may be changed by a process to any value that is less than or equal to the  
18452 hard limit. A process may (irreversibly) lower its hard limit to any value that is greater than or  
18453 equal to the soft limit. Only a process with appropriate privileges can raise a hard limit. Both  
18454 hard and soft limits can be changed in a single call to *setrlimit()* subject to the constraints  
18455 described above.18456 The value RLIM\_INFINITY, defined in *<sys/resource.h>*, shall be considered to be larger than  
18457 any other limit value. If a call to *getrlimit()* returns RLIM\_INFINITY for a resource, it means the  
18458 implementation shall not enforce limits on that resource. Specifying RLIM\_INFINITY as any  
18459 resource limit value on a successful call to *setrlimit()* inhibits enforcement of that resource limit.

18460 The following resources are defined:

|       |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 18461 | RLIMIT_CORE   | This is the maximum size of a core file in bytes that may be created by a process. A limit of 0 shall prevent the creation of a core file. If this limit is exceeded, the writing of a core file shall terminate at this size.                                                                                                                                                                                  |
| 18462 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18463 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18464 | RLIMIT_CPU    | This is the maximum amount of CPU time in seconds used by a process. If this limit is exceeded, SIGXCPU shall be generated for the process. If the process is catching or ignoring SIGXCPU, or all threads belonging to that process are blocking SIGXCPU, the behavior is unspecified.                                                                                                                         |
| 18465 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18466 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18467 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18468 | RLIMIT_DATA   | This is the maximum size of a process' data segment in bytes. If this limit is exceeded, the <i>malloc()</i> function shall fail with <i>errno</i> set to [ENOMEM].                                                                                                                                                                                                                                             |
| 18469 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18470 | RLIMIT_FSIZE  | This is the maximum size of a file in bytes that may be created by a process. If a write or truncate operation would cause this limit to be exceeded, SIGXFSZ shall be generated for the thread. If the thread is blocking, or the process is catching or ignoring SIGXFSZ, continued attempts to increase the size of a file from end-of-file to beyond the limit shall fail with <i>errno</i> set to [EFBIG]. |
| 18471 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18472 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18473 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18474 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18475 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18476 | RLIMIT_NOFILE | This is a number one greater than the maximum value that the system may assign to a newly-created descriptor. If this limit is exceeded, functions that allocate new file descriptors may fail with <i>errno</i> set to [EMFILE]. This limit constrains the number of file descriptors that a process may allocate.                                                                                             |
| 18477 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18478 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18479 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18480 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18481 | RLIMIT_STACK  | This is the maximum size of a process' stack in bytes. The implementation does not automatically grow the stack beyond this limit. If this limit is exceeded, SIGSEGV shall be generated for the thread. If the                                                                                                                                                                                                 |
| 18482 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18483 |               |                                                                                                                                                                                                                                                                                                                                                                                                                 |

18484 thread is blocking SIGSEGV, or the process is ignoring or catching  
 18485 SIGSEGV and has not made arrangements to use an alternate stack, the  
 18486 disposition of SIGSEGV shall be set to SIG\_DFL before it is generated.

18487 RLIMIT\_AS This is the maximum size of a process' total available memory, in bytes. If  
 18488 this limit is exceeded, the *malloc()* and *mmap()* functions shall fail with  
 18489 *errno* set to [ENOMEM]. In addition, the automatic stack growth fails  
 18490 with the effects outlined above.

18491 When using the *getrlimit()* function, if a resource limit can be represented correctly in an object  
 18492 of type **rlim\_t**, then its representation is returned; otherwise, if the value of the resource limit is  
 18493 equal to that of the corresponding saved hard limit, the value returned shall be  
 18494 RLIM\_SAVED\_MAX; otherwise, the value returned shall be RLIM\_SAVED\_CUR.

18495 When using the *setrlimit()* function, if the requested new limit is RLIM\_INFINITY, the new limit  
 18496 shall be “no limit”; otherwise, if the requested new limit is RLIM\_SAVED\_MAX, the new limit  
 18497 shall be the corresponding saved hard limit; otherwise, if the requested new limit is  
 18498 RLIM\_SAVED\_CUR, the new limit shall be the corresponding saved soft limit; otherwise, the  
 18499 new limit shall be the requested value. In addition, if the corresponding saved limit can be  
 18500 represented correctly in an object of type **rlim\_t** then it shall be overwritten with the new limit.

18501 The result of setting a limit to RLIM\_SAVED\_MAX or RLIM\_SAVED\_CUR is unspecified unless  
 18502 a previous call to *getrlimit()* returned that value as the soft or hard limit for the corresponding  
 18503 resource limit.

18504 The determination of whether a limit can be correctly represented in an object of type **rlim\_t** is  
 18505 implementation-defined. For example, some implementations permit a limit whose value is  
 18506 greater than RLIM\_INFINITY and others do not.

18507 The *exec* family of functions also cause resource limits to be saved.

#### 18508 RETURN VALUE

18509 Upon successful completion, *getrlimit()* and *setrlimit()* shall return 0. Otherwise, these functions  
 18510 shall return -1 and set *errno* to indicate the error.

#### 18511 ERRORS

18512 The *getrlimit()* and *setrlimit()* functions shall fail if:

18513 [EINVAL] An invalid *resource* was specified; or in a *setrlimit()* call, the new *rlim\_cur*  
 18514 exceeds the new *rlim\_max*.

18515 [EPERM] The limit specified to *setrlimit()* would have raised the maximum limit value,  
 18516 and the calling process does not have appropriate privileges.

18517 The *setrlimit()* function may fail if:

18518 [EINVAL] The limit specified cannot be lowered because current usage is already higher  
 18519 than the limit.

#### 18520 EXAMPLES

18521 None.

#### 18522 APPLICATION USAGE

18523 If a process attempts to set the hard limit or soft limit for RLIMIT\_NOFILE to less than the value  
 18524 of {\_POSIX\_OPEN\_MAX} from <**limits.h**>, unexpected behavior may occur.

18525 **RATIONALE**

18526           None.

18527 **FUTURE DIRECTIONS**

18528           None.

18529 **SEE ALSO**18530           *exec*, *fork()*, *malloc()*, *open()*, *sigaltstack()*, *sysconf()*, *ulimit()*, the Base Definitions volume of

18531           IEEE Std. 1003.1-200x, &lt;stropts.h&gt;, &lt;sys/resource.h&gt;

18532 **CHANGE HISTORY**

18533           First released in Issue 4, Version 2.

18534 **Issue 5**

18535           Moved from X/OPEN UNIX extension to BASE and an APPLICATION USAGE section is added.

18536           Large File Summit extensions are added.

18537 **NAME**

18538           getrusage — get information about resource utilization

18539 **SYNOPSIS**

18540 XSI       #include &lt;sys/resource.h&gt;

18541       int getrusage(int *who*, struct rusage \**r\_usage*);

18542

18543 **DESCRIPTION**

18544       The *getrusage()* function provides measures of the resources used by the current process or its  
18545       terminated and waited-for child processes. If the value of the *who* argument is RUSAGE\_SELF,  
18546       information shall be returned about resources used by the current process. If the value of the *who*  
18547       argument is RUSAGE\_CHILDREN, information shall be returned about resources used by the  
18548       terminated and waited-for children of the current process. If the child is never waited for (for  
18549       example, if the parent has SA\_NOCLDWAIT set or sets SIGCHLD to SIG\_IGN), the resource  
18550       information for the child process is discarded and not included in the resource information  
18551       provided by *getrusage()*.

18552       The *r\_usage* argument is a pointer to an object of type **struct rusage** in which the returned  
18553       information is stored.

18554 **RETURN VALUE**

18555       Upon successful completion, *getrusage()* shall return 0; otherwise, -1 shall be returned and *errno*  
18556       set to indicate the error.

18557 **ERRORS**18558       The *getrusage()* function shall fail if:18559       [EINVAL]       The value of the *who* argument is not valid.18560 **EXAMPLES**18561       **Using getrusage()**

18562       The following example returns information about the resources used by the current process.

18563       #include &lt;sys/resource.h&gt;

18564       ...

18565       int who = RUSAGE\_SELF;

18566       struct rusage usage;

18567       int ret;

18568       ret = getrusage(who, &amp;usage);

18569 **APPLICATION USAGE**

18570       None.

18571 **RATIONALE**

18572       None.

18573 **FUTURE DIRECTIONS**

18574       None.

18575 **SEE ALSO**18576       *exit()*, *sigaction()*, *time()*, *times()*, *wait()*, the Base Definitions volume of IEEE Std. 1003.1-200x,

18577       &lt;sys/resource.h&gt;

18578 **CHANGE HISTORY**

18579 First released in Issue 4, Version 2.

18580 **Issue 5**

18581 Moved from X/OPEN UNIX extension to BASE.

18582 **NAME**

18583       gets — get a string from a stdin stream

18584 **SYNOPSIS**

18585       #include <stdio.h>

18586       char \*gets(char \*s);

18587 **DESCRIPTION**

18588 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
18589 conflict between the requirements described here and the ISO C standard is unintentional. This  
18590 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

18591       The *gets()* function shall read bytes from the standard input stream, *stdin*, into the array pointed  
18592 to by *s*, until a newline is read or an end-of-file condition is encountered. Any newline is  
18593 discarded and a null byte is placed immediately after the last byte read into the array.

18594 **CX**       The *gets()* function may mark the *st\_atime* field of the file associated with *stream* for update. The  
18595 *st\_atime* field shall be marked for update by the first successful execution of *fgetc()*, *fgets()*,  
18596 *fread()*, *getc()*, *getchar()*, *gets()*, *fscanf()*, or *scanf()* using *stream* that returns data not supplied by  
18597 a prior call to *ungetc()*.

18598 **RETURN VALUE**

18599       Upon successful completion, *gets()* shall return *s*. If the stream is at end-of-file, the end-of-file  
18600 indicator for the stream shall be set and *gets()* shall return a null pointer. If a read error occurs,  
18601 **CX**       the error indicator for the stream shall be set, *gets()* shall return a null pointer and set *errno* to  
18602 indicate the error.

18603 **ERRORS**

18604       Refer to *fgetc()*.

18605 **EXAMPLES**

18606       None.

18607 **APPLICATION USAGE**

18608       Reading a line that overflows the array pointed to by *s* results in undefined behavior. The use of  
18609 *fgets()* is recommended.

18610       Since the user cannot specify the length of the buffer passed to *gets()*, use of this function is  
18611 discouraged. The length of the string read is unlimited. It is possible to overflow this buffer in  
18612 such a way as to cause applications to fail, or possible system security violations.

18613       It is recommended that the *fgets()* function should be used to read input lines.

18614 **RATIONALE**

18615       None.

18616 **FUTURE DIRECTIONS**

18617       None.

18618 **SEE ALSO**

18619       *feof()*, *ferror()*, *fgets()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdio.h**>

18620 **CHANGE HISTORY**

18621       First released in Issue 1. Derived from Issue 1 of the SVID.

18622 **Issue 4**

18623       In the DESCRIPTION:

- 18624       • The text is changed to make it clear that the function reads bytes rather than (possibly multi-  
18625       byte) characters.

- 18626 • The list of functions that may cause the *st\_atime* field to be updated is revised.
- 18627 **Issue 6**
- 18628 Extensions beyond the ISO C standard are now marked.

18629 **NAME**

18630        getservbyname — network services database functions

18631 **SYNOPSIS**

18632        #include <netdb.h>

18633        struct servent \*getservbyname(const char \*name, const char \*proto);

18634 **DESCRIPTION**

18635        Refer to *endservent()*.

18636 **NAME**

18637        getservbyport — network services database functions

18638 **SYNOPSIS**

18639        #include &lt;netdb.h&gt;

18640        struct servent \*getservbyport(int *port*, const char \**proto*);18641 **DESCRIPTION**18642        Refer to *endservent()*.

18643 **NAME**

18644        getservent — network services database functions

18645 **SYNOPSIS**

18646        #include <netdb.h>

18647        struct servent \*getservent(void);

18648 **DESCRIPTION**

18649        Refer to *endservent()*.

18650 **NAME**

18651 getsid — get the process group ID of a session leader

18652 **SYNOPSIS**

18653 XSI #include &lt;unistd.h&gt;

18654 pid\_t getsid(pid\_t pid);

18655

18656 **DESCRIPTION**18657 The *getsid()* function obtains the process group ID of the process that is the session leader of the  
18658 process specified by *pid*. If *pid* is (**pid\_t**)0, it specifies the calling process.18659 **RETURN VALUE**18660 Upon successful completion, *getsid()* shall return the process group ID of the session leader of  
18661 the specified process. Otherwise, it shall return (**pid\_t**)-1 and set *errno* to indicate the error.18662 **ERRORS**18663 The *getsid()* function shall fail if:18664 [EPERM] The process specified by *pid* is not in the same session as the calling process,  
18665 and the implementation does not allow access to the process group ID of the  
18666 session leader of that process from the calling process.18667 [ESRCH] There is no process with a process ID equal to *pid*.18668 **EXAMPLES**

18669 None.

18670 **APPLICATION USAGE**

18671 None.

18672 **RATIONALE**

18673 None.

18674 **FUTURE DIRECTIONS**

18675 None.

18676 **SEE ALSO**18677 *exec*, *fork()*, *getpid()*, *getpgid()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
18678 IEEE Std. 1003.1-200x, <unistd.h>18679 **CHANGE HISTORY**

18680 First released in Issue 4, Version 2.

18681 **Issue 5**

18682 Moved from X/OPEN UNIX extension to BASE.

18683 **NAME**

18684 getsockname — get the socket name

18685 **SYNOPSIS**

18686 #include &lt;sys/socket.h&gt;

18687 int getsockname(int *socket*, struct sockaddr \*restrict *address*,  
18688 socklen\_t \**address\_len*);18689 **DESCRIPTION**18690 The *getsockname()* function shall retrieve the locally-bound name of the specified socket, store  
18691 this address in the **sockaddr** structure pointed to by the *address* argument, and store the length of  
18692 this address in the object pointed to by the *address\_len* argument.18693 If the actual length of the address is greater than the length of the supplied **sockaddr** structure,  
18694 the stored address shall be truncated.18695 If the socket has not been bound to a local name, the value stored in the object pointed to by  
18696 *address* is unspecified.18697 **RETURN VALUE**18698 Upon successful completion, 0 shall be returned, the *address* argument shall point to the address  
18699 of the socket, and the *address\_len* argument shall point to the length of the address. Otherwise, -1  
18700 shall be returned and *errno* set to indicate the error.18701 **ERRORS**18702 The *getsockname()* function shall fail if:18703 [EBADF] The *socket* argument is not a valid file descriptor.18704 [ENOTSOCK] The *socket* argument does not refer to a socket.

18705 [EOPNOTSUPP] The operation is not supported for this socket's protocol.

18706 The *getsockname()* function may fail if:

18707 [EINVAL] The socket has been shut down.

18708 [ENOBUFS] Insufficient resources were available in the system to complete the function.

18709 **EXAMPLES**

18710 None.

18711 **APPLICATION USAGE**

18712 None.

18713 **RATIONALE**

18714 None.

18715 **FUTURE DIRECTIONS**

18716 None.

18717 **SEE ALSO**18718 *accept()*, *bind()*, *getpeername()*, *socket()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
18719 <sys/socket.h>18720 **CHANGE HISTORY**

18721 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

18722 The **restrict** keyword is added to the *getsockname()* prototype for alignment with the  
18723 ISO/IEC 9899:1999 standard.

18724 **NAME**

18725 getsockopt — get the socket options

18726 **SYNOPSIS**

18727 #include &lt;sys/socket.h&gt;

```
18728 int getsockopt(int socket, int level, int option_name,
18729 void *restrict option_value, socklen_t *restrict option_len);
```

18730 **DESCRIPTION**18731 The *getsockopt()* function manipulates options associated with a socket.

18732 The *getsockopt()* function shall retrieve the value for the option specified by the *option\_name*  
 18733 argument for the socket specified by the *socket* argument. If the size of the option value is greater  
 18734 than *option\_len*, the value stored in the object pointed to by the *option\_value* argument shall be  
 18735 silently truncated. Otherwise, the object pointed to by the *option\_len* argument shall be modified  
 18736 to indicate the actual length of the value.

18737 The *level* argument specifies the protocol level at which the option resides. To retrieve options at  
 18738 the socket level, specify the *level* argument as SOL\_SOCKET. To retrieve options at other levels,  
 18739 supply the appropriate level identifier for the protocol controlling the option. For example, to  
 18740 indicate that an option is interpreted by the TCP (Transmission Control Protocol), set *level* to  
 18741 IPPROTO\_TCP as defined in the <netinet/in.h> header.

18742 The socket in use may require the process to have appropriate privileges to use the *getsockopt()*  
 18743 function.

18744 The *option\_name* argument specifies a single option to be retrieved. It can be one of the following  
 18745 values defined in <sys/socket.h>:

18746 SO\_DEBUG Reports whether debugging information is being recorded. This option  
 18747 stores an **int** value. This is a Boolean option.

18748 SO\_ACCEPTCONN Reports whether socket listening is enabled. This option stores an **int**  
 18749 value. This is a Boolean option.

18750 SO\_BROADCAST Reports whether transmission of broadcast messages is supported, if this  
 18751 is supported by the protocol. This option stores an **int** value. This is a  
 18752 Boolean option.

18753 SO\_REUSEADDR Reports whether the rules used in validating addresses supplied to *bind()*  
 18754 should allow reuse of local addresses, if this is supported by the protocol.  
 18755 This option stores an **int** value. This is a Boolean option.

18756 SO\_KEEPALIVE Reports whether connections are kept active with periodic transmission  
 18757 of messages, if this is supported by the protocol.

18758 If the connected socket fails to respond to these messages, the connection  
 18759 is broken and threads writing to that socket are notified with a SIGPIPE  
 18760 signal. This option stores an **int** value.

18761 This is a Boolean option.

18762 SO\_LINGER Reports whether the socket lingers on *close()* if data is present. If  
 18763 SO\_LINGER is set, the system blocks the process during *close()* until it  
 18764 can transmit the data or until the end of the interval indicated by the  
 18765 *l\_linger* member, whichever comes first. If SO\_LINGER is not specified,  
 18766 and *close()* is issued, the system handles the call in a way that allows the  
 18767 process to continue as quickly as possible. This option stores a **linger**  
 18768 structure.

|       |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 18769 | SO_OOBINLINE | Reports whether the socket leaves received out-of-band data (data marked urgent) inline. This option stores an <b>int</b> value. This is a Boolean option.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 18770 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18771 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18772 | SO_SNDBUF    | Reports send buffer size information. This option stores an <b>int</b> value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18773 | SO_RCVBUF    | Reports receive buffer size information. This option stores an <b>int</b> value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 18774 | SO_ERROR     | Reports information about error status and clears it. This option stores an <b>int</b> value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 18775 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18776 | SO_TYPE      | Reports the socket type. This option stores an <b>int</b> value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 18777 | SO_DONTROUTE | Reports whether outgoing messages bypass the standard routing facilities. The destination shall be on a directly-connected network, and messages are directed to the appropriate network interface according to the destination address. The effect, if any, of this option depends on what protocol is in use. This option stores an <b>int</b> value. This is a Boolean option.                                                                                                                                                                                                                             |
| 18778 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18779 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18780 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18781 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18782 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18783 | SO_RCVLOWAT  | Reports the minimum number of bytes to process for socket input operations. The default value for SO_RCVLOWAT is 1. If SO_RCVLOWAT is set to a larger value, blocking receive calls normally wait until they have received the smaller of the low water mark value or the requested amount. (They may return less than the low water mark if an error occurs, a signal is caught, or the type of data next in the receive queue is different than that returned; for example, out-of-band data.) This option stores an <b>int</b> value. Note that not all implementations allow this option to be retrieved. |
| 18784 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18785 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18786 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18787 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18788 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18789 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18790 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18791 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18792 | SO_RCVTIMEO  | Reports the timeout value for input operations. This option stores a <b>timeval</b> structure with the number of seconds and microseconds specifying the limit on how long to wait for an input operation to complete. If a receive operation has blocked for this much time without receiving additional data, it shall return with a partial count or <i>errno</i> set to [EAGAIN] or [EWOULDBLOCK] if no data was received. The default for this option is zero, which indicates that a receive operation shall not time out. Note that not all implementations allow this option to be retrieved.         |
| 18793 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18794 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18795 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18796 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18797 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18798 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18799 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18800 | SO_SNDLOWAT  | Reports the minimum number of bytes to process for socket output operations. Non-blocking output operations shall process no data if flow control does not allow the smaller of the send low water mark value or the entire request to be processed. This option stores an <b>int</b> value. Note that not all implementations allow this option to be retrieved.                                                                                                                                                                                                                                             |
| 18801 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18802 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18803 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18804 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18805 | SO_SNDTIMEO  | Reports the timeout value specifying the amount of time that an output function blocks because flow control prevents data from being sent. If a send operation has blocked for this time, it shall return with a partial count or with <i>errno</i> set to [EAGAIN] or [EWOULDBLOCK] if no data were sent. The default for this option is zero, which indicates that a send operation shall not time out. The option stores a <b>timeval</b> structure. Note that not all implementations allow this option to be retrieved.                                                                                  |
| 18806 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18807 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18808 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18809 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18810 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18811 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18812 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18813 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 18814 |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|       |              | For Boolean options, a zero value indicates that the option is disabled and a non-zero value indicates that the option is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|       |              | Options at other protocol levels vary in format and name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

18815 The socket in use may require the process to have appropriate privileges to use the *getsockopt()*  
18816 function.

#### 18817 RETURN VALUE

18818 Upon successful completion, *getsockopt()* shall return 0; otherwise, -1 shall be returned and *errno*  
18819 set to indicate the error.

#### 18820 ERRORS

18821 The *getsockopt()* function shall fail if:

18822 [EBADF] The *socket* argument is not a valid file descriptor.

18823 [EINVAL] The specified option is invalid at the specified socket level.

18824 [ENOPROTOOPT]

18825 The option is not supported by the protocol.

18826 [ENOTSOCK] The *socket* argument does not refer to a socket.

18827 The *getsockopt()* function may fail if:

18828 [EACCES] The calling process does not have the appropriate privileges.

18829 [EINVAL] The socket has been shut down.

18830 [ENOBUFS] Insufficient resources are available in the system to complete the function.

#### 18831 EXAMPLES

18832 None.

#### 18833 APPLICATION USAGE

18834 None.

#### 18835 RATIONALE

18836 None.

#### 18837 FUTURE DIRECTIONS

18838 None.

#### 18839 SEE ALSO

18840 *bind()*, *close()*, *endprotoent()*, *setsockopt()*, *socket()*, the Base Definitions volume of  
18841 IEEE Std. 1003.1-200x, <sys/socket.h>, <netinet/in.h>

#### 18842 CHANGE HISTORY

18843 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

18844 The **restrict** keyword is added to the *getsockopt()* prototype for alignment with the  
18845 ISO/IEC 9899:1999 standard.

18846 **NAME**

18847       getsubopt — parse suboption arguments from a string

18848 **SYNOPSIS**

18849 XSI       #include &lt;stdlib.h&gt;

18850       int getsubopt(char \*\*optionp, char \* const \*tokens, char \*\*valuep);

18851

18852 **DESCRIPTION**

18853       The *getsubopt()* function parses suboption arguments in a flag argument that was initially parsed  
 18854       by *getopt()*. The application shall ensure that these suboption arguments are separated by  
 18855       commas and may consist of either a single token, or a token-value pair separated by an equal  
 18856       sign. Because commas delimit suboption arguments in the option string, they are not allowed to  
 18857       be part of the suboption arguments or the value of a suboption argument. Similarly, because the  
 18858       equal sign separates a token from its value, undefined behavior will result if the application  
 18859       passes a token that contains an equal sign.

18860       The *getsubopt()* function takes the address of a pointer to the option argument string, a vector of  
 18861       possible tokens, and the address of a value string pointer. If the option argument string at  
 18862       \**optionp* contains only one suboption argument, *getsubopt()* updates \**optionp* to point to the null  
 18863       at the end of the string. Otherwise, it isolates the suboption argument by replacing the comma  
 18864       separator with a null, and updates \**optionp* to point to the start of the next suboption argument.  
 18865       If the suboption argument has an associated value, *getsubopt()* updates \**valuep* to point to the  
 18866       value's first character. Otherwise, it sets \**valuep* to a null pointer.

18867       The token vector is organized as a series of pointers to strings. The end of the token vector is  
 18868       identified by a null pointer.

18869       When *getsubopt()* returns, if \**valuep* is not a null pointer, then the suboption argument processed  
 18870       included a value. The calling program may use this information to determine whether the  
 18871       presence or lack of a value for this suboption is an error.

18872       Additionally, when *getsubopt()* fails to match the suboption argument with the tokens in the  
 18873       *tokens* array, the calling program should decide if this is an error, or if the unrecognized option  
 18874       should be passed on to another program.

18875 **RETURN VALUE**

18876       The *getsubopt()* function shall return the index of the matched token string, or -1 if no token  
 18877       strings were matched.

18878 **ERRORS**

18879       No errors are defined.

18880 **EXAMPLES**

```
18881 #include <stdio.h>
18882 #include <stdlib.h>

18883 int do_all;
18884 const char *type;
18885 int read_size;
18886 int write_size;
18887 int read_only;

18888 enum
18889 {
18890 RO_OPTION = 0,
18891 RW_OPTION,
```

```

18892 READ_SIZE_OPTION,
18893 WRITE_SIZE_OPTION
18894 };
18895 const char *mount_opts[] =
18896 {
18897 [RO_OPTION] = "ro",
18898 [RW_OPTION] = "rw",
18899 [READ_SIZE_OPTION] = "rsize",
18900 [WRITE_SIZE_OPTION] = "wsize",
18901 NULL
18902 };
18903 int
18904 main(int argc, char *argv[])
18905 {
18906 char *subopts, *value;
18907 int opt;
18908
18909 while ((opt = getopt(argc, argv, "at:o:")) != -1)
18910 switch(opt)
18911 {
18912 case 'a':
18913 do_all = 1;
18914 break;
18915 case 't':
18916 type = optarg;
18917 break;
18918 case 'o':
18919 subopts = optarg;
18920 while (*subopts != ' ')
18921 switch(getsubopt(&subopts, mount_opts, &value))
18922 {
18923 case RO_OPTION:
18924 read_only = 1;
18925 break;
18926 case RW_OPTION:
18927 read_only = 0;
18928 break;
18929 case READ_SIZE_OPTION:
18930 if (value == NULL)
18931 abort();
18932 read_size = atoi(value);
18933 break;
18934 case WRITE_SIZE_OPTION:
18935 if (value == NULL)
18936 abort();
18937 write_size = atoi(value);
18938 break;
18939 default:
18940 /* Unknown suboption. */
18941 printf("Unknown suboption '%s'\n", value);
18942 break;
18943 }
18944 }

```

```

18943 break;
18944 default:
18945 abort();
18946 }
18947 /* Do the real work. */
18948 return 0;
18949 }

```

### 18950 Parsing Suboptions

18951 The following example uses the *getsubopt()* function to parse a value argument in the *optarg*  
 18952 external variable returned by a call to *getopt()*.

```

18953 #include <stdlib.h>
18954 ...
18955 char *tokens[] = {"HOME", "PATH", "LOGNAME", (char *) NULL };
18956 char *value;
18957 int opt, index;
18958 while ((opt = getopt(argc, argv, "e:")) != -1) {
18959 switch(opt) {
18960 case 'e' :
18961 while ((index = getsubopt(&optarg, tokens, &value)) != -1) {
18962 switch(index) {
18963 ...
18964 }
18965 break;
18966 ...
18967 }
18968 }
18969 ...

```

### 18970 APPLICATION USAGE

18971 None.

### 18972 RATIONALE

18973 None.

### 18974 FUTURE DIRECTIONS

18975 None.

### 18976 SEE ALSO

18977 *getopt()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

### 18978 CHANGE HISTORY

18979 First released in Issue 4, Version 2.

### 18980 Issue 5

18981 Moved from X/OPEN UNIX extension to BASE.

18982 **NAME**

18983           gettimeofday — get the date and time

18984 **SYNOPSIS**

18985 XSI       #include &lt;sys/time.h&gt;

18986       int gettimeofday(struct timeval \*restrict *tp*, void \*tzp);

18987

18988 **DESCRIPTION**

18989       The *gettimeofday()* function obtains the current time, expressed as seconds and microseconds since the Epoch, and stores it in the **timeval** structure pointed to by *tp*. The resolution of the system clock is unspecified.

18992       If *tzp* is not a null pointer, the behavior is unspecified.

18993 **RETURN VALUE**

18994       The *gettimeofday()* function shall return 0 and no value shall be reserved to indicate an error.

18995 **ERRORS**

18996       No errors are defined.

18997 **EXAMPLES**

18998       None.

18999 **APPLICATION USAGE**

19000       None.

19001 **RATIONALE**

19002       None.

19003 **FUTURE DIRECTIONS**

19004       None.

19005 **SEE ALSO**

19006       *ctime()*, *ftime()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/time.h>

19007 **CHANGE HISTORY**

19008       First released in Issue 4, Version 2.

19009 **Issue 5**

19010       Moved from X/OPEN UNIX extension to BASE.

19011 **Issue 6**

19012       The DESCRIPTION is updated to refer to “seconds since the Epoch” rather than “seconds since 00:00:00 UTC (Coordinated Universal Time), January 1 1970” for consistency with other *time* functions.

19015 **NAME**

19016       getuid — get a real user ID

19017 **SYNOPSIS**

19018       #include &lt;unistd.h&gt;

19019       uid\_t getuid(void);

19020 **DESCRIPTION**19021       The *getuid()* function shall return the real user ID of the calling process.19022 **RETURN VALUE**19023       The *getuid()* function is always successful and no return value is reserved to indicate the error.19024 **ERRORS**

19025       No errors are defined.

19026 **EXAMPLES**19027       **Setting the Effective User ID to the Real User ID**19028       The following example sets the effective user ID and the real user ID of the current process to the  
19029       real user ID of the caller.19030       #include <unistd.h>  
19031       #include <sys/types.h>  
19032       ...  
19033       setreuid(getuid(), getuid());  
19034       ...19035 **APPLICATION USAGE**

19036       None.

19037 **RATIONALE**

19038       None.

19039 **FUTURE DIRECTIONS**

19040       None.

19041 **SEE ALSO**19042       *getegid()*, *geteuid()*, *getgid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the Base  
19043       Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>19044 **CHANGE HISTORY**

19045       First released in Issue 1. Derived from Issue 1 of the SVID.

19046 **Issue 4**19047       The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
19048       XSI-conformant systems.

19049       The &lt;unistd.h&gt; header is added to the SYNOPSIS section.

19050       The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 19051
- The argument list is explicitly defined as **void**.

19052 **Issue 6**

19053       In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

19054       The following new requirements on POSIX implementations derive from alignment with the  
19055       Single UNIX Specification:

19056  
19057  
19058

- The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was required for conforming implementations of previous POSIX specifications, it was not required for UNIX applications.

19059 **NAME**

19060           getutxent, getutxid, getutxline — get user accounting database entries

19061 **SYNOPSIS**

19062 XSI       #include <utmpx.h>

19063           struct utmpx \*getutxent(void);

19064           struct utmpx \*getutxid(const struct utmpx \*id);

19065           struct utmpx \*getutxline(const struct utmpx \*line);

19066

19067 **DESCRIPTION**

19068           Refer to *endutxent()*.

19069 **NAME**19070 `getwc` — get a wide character from a stream19071 **SYNOPSIS**19072 `#include <stdio.h>`19073 `#include <wchar.h>`19074 `wint_t getwc(FILE *stream);`19075 **DESCRIPTION**

19076 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
19077 conflict between the requirements described here and the ISO C standard is unintentional. This  
19078 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

19079 The `getwc()` function is equivalent to `fgetwc()`, except that if it is implemented as a macro it may  
19080 evaluate `stream` more than once, so the argument should never be an expression with side effects.

19081 **RETURN VALUE**19082 Refer to `fgetwc()`.19083 **ERRORS**19084 Refer to `fgetwc()`.19085 **EXAMPLES**

19086 None.

19087 **APPLICATION USAGE**

19088 Because it may be implemented as a macro, `getwc()` may treat incorrectly a `stream` argument with  
19089 side effects. In particular, `getwc(*f++)` does not necessarily work as expected. Therefore, use of  
19090 this function is not recommended; `fgetwc()` should be used instead.

19091 **RATIONALE**

19092 None.

19093 **FUTURE DIRECTIONS**

19094 None.

19095 **SEE ALSO**19096 `fgetwc()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdio.h>`, `<wchar.h>`19097 **CHANGE HISTORY**

19098 First released as a World-wide Portability Interface in Issue 4. Derived from the MSE working  
19099 draft.

19100 **Issue 5**19101 The Optional Header (OH) marking is removed from `<stdio.h>`.

19102 **NAME**

19103       getwchar — get a wide character from a stdin stream

19104 **SYNOPSIS**

19105       #include <wchar.h>

19106       wint\_t getwchar(void);

19107 **DESCRIPTION**

19108 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
19109       conflict between the requirements described here and the ISO C standard is unintentional. This  
19110       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

19111       The *getwchar()* function is equivalent to *getwc(stdin)*.

19112 **RETURN VALUE**

19113       Refer to *fgetwc()*.

19114 **ERRORS**

19115       Refer to *fgetwc()*.

19116 **EXAMPLES**

19117       None.

19118 **APPLICATION USAGE**

19119       If the **wint\_t** value returned by *getwchar()* is stored into a variable of type **wchar\_t** and then  
19120       compared against the **wint\_t** macro WEOF, the result may be incorrect. Only the **wint\_t** type is  
19121       guaranteed to be able to represent any wide character and WEOF.

19122 **RATIONALE**

19123       None.

19124 **FUTURE DIRECTIONS**

19125       None.

19126 **SEE ALSO**

19127       *fgetwc()*, *getwc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>

19128 **CHANGE HISTORY**

19129       First released as a World-wide Portability Interface in Issue 4. Derived from the MSE working  
19130       draft.

19131 **NAME**

19132           getwd — get the current working directory path name (**LEGACY**)

19133 **SYNOPSIS**

19134 XSI       #include <unistd.h>

19135       char \*getwd(char \*path\_name);

19136

19137 **DESCRIPTION**

19138       The *getwd()* function determines an absolute path name of the current working directory of the  
19139       calling process, and copies that path name into the array pointed to by the *path\_name* argument.

19140       If the length of the path name of the current working directory is greater than ({PATH\_MAX}+1)  
19141       including the null byte, *getwd()* shall fail and return a null pointer.

19142 **RETURN VALUE**

19143       Upon successful completion, a pointer to the string containing the absolute path name of the  
19144       current working directory shall be returned. Otherwise, *getwd()* shall return a null pointer and  
19145       the contents of the array pointed to by *path\_name* are undefined.

19146 **ERRORS**

19147       No errors are defined.

19148 **EXAMPLES**

19149       None.

19150 **APPLICATION USAGE**

19151       For applications portability, the *getcwd()* function should be used to determine the current  
19152       working directory instead of *getwd()*.

19153 **RATIONALE**

19154       Since the user cannot specify the length of the buffer passed to *getwd()*, use of this function is  
19155       discouraged. The length of a path name described in {PATH\_MAX} is file system-dependent and  
19156       may vary from one mount point to another, or might even be unlimited. It is possible to  
19157       overflow this buffer in such a way as to cause applications to fail, or possible system security  
19158       violations.

19159       It is recommended that the *getcwd()* function should be used to determine the current working  
19160       directory.

19161 **FUTURE DIRECTIONS**

19162       This function may be withdrawn in a future version.

19163 **SEE ALSO**

19164       *getcwd()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <unistd.h>

19165 **CHANGE HISTORY**

19166       First released in Issue 4, Version 2.

19167 **Issue 5**

19168       Moved from X/OPEN UNIX extension to BASE.

19169 **Issue 6**

19170       This function is marked LEGACY.

## 19171 NAME

19172 glob, globfree — generate path names matching a pattern

## 19173 SYNOPSIS

19174 #include <glob.h>

```
19175 int glob(const char *restrict pattern, int flags,
19176 int(*restrict errfunc)(const char *retrict epath, int eerrno),
19177 glob_t *restrict pglob);
19178 void globfree(glob_t *pglob);
```

## 19179 DESCRIPTION

19180 The *glob()* function is a path name generator that implements the rules defined in the Shell and  
 19181 Utilities volume of IEEE Std. 1003.1-200x, Section 2.14, Pattern Matching Notation, with optional  
 19182 support for rule 3 in the Shell and Utilities volume of IEEE Std. 1003.1-200x, Section 2.14.3,  
 19183 Patterns Used for File Name Expansion.

19184 The structure type **glob\_t** is defined in <glob.h> and includes at least the following members:

19185

19186

19187

19188

19189

| Member Type | Member Name     | Description                                            |
|-------------|-----------------|--------------------------------------------------------|
| size_t      | <i>gl_pathc</i> | Count of paths matched by <i>pattern</i> .             |
| char **     | <i>gl_pathv</i> | Pointer to a list of matched path names.               |
| size_t      | <i>gl_offs</i>  | Slots to reserve at the beginning of <i>gl_pathv</i> . |

19190 The argument *pattern* is a pointer to a path name pattern to be expanded. The *glob()* function  
 19191 matches all accessible path names against this pattern and develops a list of all path names that  
 19192 match. In order to have access to a path name, *glob()* requires search permission on every  
 19193 component of a path except the last, and read permission on each directory of any file name  
 19194 component of *pattern* that contains any of the following special characters: '\*', '?', and '['.

19195 The *glob()* function stores the number of matched path names into *pglob->gl\_pathc* and a pointer  
 19196 to a list of pointers to path names into *pglob->gl\_pathv*. The path names are in sort order as  
 19197 defined by the current setting of the *LC\_COLLATE* category; see the Base Definitions volume of  
 19198 IEEE Std. 1003.1-200x, Section 7.3.2, *LC\_COLLATE*. The first pointer after the last path name is a  
 19199 null pointer. If the pattern does not match any path names, the returned number of matched  
 19200 paths is set to 0, and the contents of *pglob->gl\_pathv* are implementation-defined.

19201 It is the caller's responsibility to create the structure pointed to by *pglob*. The *glob()* function  
 19202 allocates other space as needed, including the memory pointed to by *gl\_pathv*. The *globfree()*  
 19203 function frees any space associated with *pglob* from a previous call to *glob()*.

19204 The *flags* argument is used to control the behavior of *glob()*. The value of *flags* is a bitwise-  
 19205 inclusive OR of zero or more of the following constants, which are defined in <glob.h>:

19206 **GLOB\_APPEND** Append path names generated to the ones from a previous call to *glob()*.

19207 **GLOB\_DOOFFS** Make use of *pglob->gl\_offs*. If this flag is set, *pglob->gl\_offs* is used to  
 19208 specify how many null pointers to add to the beginning of *pglob->gl\_pathv*. In other words,  
 19209 *pglob->gl\_pathv* shall point to *pglob->gl\_offs* null pointers, followed by *pglob->gl\_pathc*  
 19210 path name pointers, followed by a null pointer.  
 19211

19212 **GLOB\_ERR** Causes *glob()* to return when it encounters a directory that it cannot open  
 19213 or read. Ordinarily, *glob()* continues to find matches.

19214 **GLOB\_MARK** Each path name that is a directory that matches *pattern* has a slash  
 19215 appended.

|       |               |                                                                                                                                                                                                                                                                                                          |
|-------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 19216 | GLOB_NOCHECK  | Supports rule 3 in the Shell and Utilities volume of IEEE Std. 1003.1-200x, Section 2.14.3, Patterns Used for File Name Expansion. If <i>pattern</i> does not match any path name, then <i>glob()</i> shall return a list consisting of only <i>pattern</i> , and the number of matched path names is 1. |
| 19217 |               |                                                                                                                                                                                                                                                                                                          |
| 19218 |               |                                                                                                                                                                                                                                                                                                          |
| 19219 |               |                                                                                                                                                                                                                                                                                                          |
| 19220 | GLOB_NOESCAPE | Disable backslash escaping.                                                                                                                                                                                                                                                                              |
| 19221 | GLOB_NOSORT   | Ordinarily, <i>glob()</i> sorts the matching path names according to the current setting of the <i>LC_COLLATE</i> category, see the Base Definitions volume of IEEE Std. 1003.1-200x, Section 7.3.2, <i>LC_COLLATE</i> . When this flag is used, the order of path names returned is unspecified.        |
| 19222 |               |                                                                                                                                                                                                                                                                                                          |
| 19223 |               |                                                                                                                                                                                                                                                                                                          |
| 19224 |               |                                                                                                                                                                                                                                                                                                          |

The GLOB\_APPEND flag can be used to append a new set of path names to those found in a previous call to *glob()*. The following rules apply to applications when two or more calls to *glob()* are made with the same value of *pglob* and without intervening calls to *globfree()*:

- 19228 1. The first such call shall not set GLOB\_APPEND. All subsequent calls shall set it.
- 19229 2. All the calls shall set GLOB\_DOOFFS, or all shall not set it.
- 19230 3. After the second call, *pglob->gl\_pathv* points to a list containing the following:
  - 19231 a. Zero or more null pointers, as specified by GLOB\_DOOFFS and *pglob->gl\_offs*.
  - 19232 b. Pointers to the path names that were in the *pglob->gl\_pathv* list before the call, in the same order as before.
  - 19233 c. Pointers to the new path names generated by the second call, in the specified order.
- 19234 4. The count returned in *pglob->gl\_pathc* shall be the total number of path names from the two calls.
- 19235 5. The application can change any of the fields after a call to *glob()*. If it does, the application shall reset them to the original value before a subsequent call, using the same *pglob* value, to *globfree()* or *glob()* with the GLOB\_APPEND flag.

19240 If, during the search, a directory is encountered that cannot be opened or read and *errfunc* is not  
 19241 a null pointer, *glob()* calls (*\*errfunc()*) with two arguments:

- 19242 1. The *epath* argument is a pointer to the path that failed.
- 19243 2. The *errno* argument is the value of *errno* from the failure, as set by *opendir()*, *readdir()*, or  
 19244 *stat()*. (Other values may be used to report other errors not explicitly documented for  
 19245 those functions.)

19246 The following constants are defined as error return values for *glob()*:

|       |              |                                                                                           |
|-------|--------------|-------------------------------------------------------------------------------------------|
| 19247 | GLOB_ABORTED | The scan was stopped because GLOB_ERR was set or ( <i>*errfunc()</i> ) returned non-zero. |
| 19248 |              |                                                                                           |
| 19249 | GLOB_NOMATCH | The pattern does not match any existing path name, and GLOB_NOCHECK was not set in flags. |
| 19250 |              |                                                                                           |
| 19251 | GLOB_NOSPACE | An attempt to allocate memory failed.                                                     |

19252 If (*\*errfunc()*) is called and returns non-zero, or if the GLOB\_ERR flag is set in *flags*, *glob()* shall  
 19253 stop the scan and return GLOB\_ABORTED after setting *gl\_pathc* and *gl\_pathv* in *pglob* to reflect  
 19254 the paths already scanned. If GLOB\_ERR is not set and either *errfunc* is a null pointer or  
 19255 (*\*errfunc()*) returns 0, the error shall be ignored.

19256 **RETURN VALUE**

19257 Upon successful completion, *glob()* shall return 0. The argument *pglob->gl\_pathc* shall return the  
 19258 number of matched path names and the argument *pglob->gl\_pathv* shall contain a pointer to a  
 19259 null-terminated list of matched and sorted path names. However, if *pglob->gl\_pathc* is 0, the  
 19260 content of *pglob->gl\_pathv* is undefined.

19261 The *globfree()* function shall return no value.

19262 If *glob()* terminates due to an error, it shall return one of the non-zero constants defined in  
 19263 `<glob.h>`. The arguments *pglob->gl\_pathc* and *pglob->gl\_pathv* are still set as defined above.

19264 **ERRORS**

19265 No errors are defined.

19266 **EXAMPLES**

19267 One use of the GLOB\_DOOFFS flag is by applications that build an argument list for use with  
 19268 *execv()*, *execve()*, or *execvp()*. Suppose, for example, that an application wants to do the  
 19269 equivalent of:

```
19270 ls -l *.c
```

19271 but for some reason:

```
19272 system("ls -l *.c")
```

19273 is not acceptable. The application could obtain approximately the same result using the  
 19274 sequence:

```
19275 globbuf.gl_offs = 2;

 19276 glob("*.c", GLOB_DOOFFS, NULL, &globbuf);

 19277 globbuf.gl_pathv[0] = "ls";

 19278 globbuf.gl_pathv[1] = "-l";

 19279 execvp("ls", &globbuf.gl_pathv[0]);
```

19280 Using the same example:

```
19281 ls -l *.c *.h
```

19282 could be approximately simulated using GLOB\_APPEND as follows:

```
19283 globbuf.gl_offs = 2;

 19284 glob("*.c", GLOB_DOOFFS, NULL, &globbuf);

 19285 glob("*.h", GLOB_DOOFFS|GLOB_APPEND, NULL, &globbuf);

 19286 ...
```

19287 **APPLICATION USAGE**

19288 This function is not provided for the purpose of enabling utilities to perform path name  
 19289 expansion on their arguments, as this operation is performed by the shell, and utilities are  
 19290 explicitly not expected to redo this. Instead, it is provided for applications that need to do path  
 19291 name expansion on strings obtained from other sources, such as a pattern typed by a user or  
 19292 read from a file.

19293 If a utility needs to see if a path name matches a given pattern, it can use *fnmatch()*.

19294 Note that *gl\_pathc* and *gl\_pathv* have meaning even if *glob()* fails. This allows *glob()* to report  
 19295 partial results in the event of an error. However, if *gl\_pathc* is 0, *gl\_pathv* is unspecified even if  
 19296 *glob()* did not return an error.

19297 The GLOB\_NOCHECK option could be used when an application wants to expand a path name  
 19298 if wildcards are specified, but wants to treat the pattern as just a string otherwise. The *sh* utility  
 19299 might use this for option-arguments, for example.

19300 The new path names generated by a subsequent call with GLOB\_APPEND are not sorted  
 19301 together with the previous path names. This mirrors the way that the shell handles path name  
 19302 expansion when multiple expansions are done on a command line.

19303 Applications that need tilde and parameter expansion should use *wordexp()*.

#### 19304 RATIONALE

19305 It was claimed that the GLOB\_DOOFFS flag is unnecessary because it could be simulated using:

```
19306 new = (char **)malloc((n + pglob->gl_pathc + 1)
19307 * sizeof(char *));
19308 (void) memcpy(new+n, pglob->gl_pathv,
19309 pglob->gl_pathc * sizeof(char *));
19310 (void) memset(new, 0, n * sizeof(char *));
19311 free(pglob->gl_pathv);
19312 pglob->gl_pathv = new;
```

19313 However, this assumes that the memory pointed to by *gl\_pathv* is a block that was separately  
 19314 created using *malloc()*. This is not necessarily the case. An application should make no  
 19315 assumptions about how the memory referenced by fields in *pglob* was allocated. It might have  
 19316 been obtained from *malloc()* in a large chunk and then carved up within *glob()*, or it might have  
 19317 been created using a different memory allocator. It is not the intent of the standard developers to  
 19318 specify or imply how the memory used by *glob()* is managed.

19319 The GLOB\_APPEND flag would be used when an application wants to expand several different  
 19320 patterns into a single list.

#### 19321 FUTURE DIRECTIONS

19322 None.

#### 19323 SEE ALSO

19324 *exec*, *fnmatch()*, *opendir()*, *readdir()*, *stat()*, *wordexp()*, the Base Definitions volume of  
 19325 IEEE Std. 1003.1-200x, <*glob.h*>, the Shell and Utilities volume of IEEE Std. 1003.1-200x

#### 19326 CHANGE HISTORY

19327 First released in Issue 4. Derived from the ISO POSIX-2 standard.

#### 19328 Issue 5

19329 Moved from POSIX2 C-language Binding to BASE.

#### 19330 Issue 6

19331 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

19332 The **restrict** keyword is added to the *glob()* prototype for alignment with the ISO/IEC 9899:1999  
 19333 standard.

## 19334 NAME

19335 gmtime, gmtime\_r — convert a time value to a broken-down UTC time

## 19336 SYNOPSIS

19337 #include <time.h>

19338 struct tm \*gmtime(const time\_t \*timer);

19339 TSF struct tm \*gmtime\_r(const time\_t \*restrict timer, struct tm \*restrict result);

19340

## 19341 DESCRIPTION

19342 CX For *gmtime()*: The functionality described on this reference page is aligned with the ISO C standard. Any conflict between the requirements described here and the ISO C standard is unintentional. This volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

19345 The *gmtime()* function shall convert the time in seconds since the Epoch pointed to by *timer* into a broken-down time, expressed as Coordinated Universal Time (UTC).

19347 CX The *gmtime()* function need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

19349 TSF The *gmtime\_r()* function shall convert the time in seconds since the Epoch pointed to by *timer* into a broken-down time expressed as Coordinated Universal Time (UTC). The broken-down time is stored in the structure referred to by *result*. The *gmtime\_r()* function shall also return the address of the same structure.

## 19353 RETURN VALUE

19354 The *gmtime()* function shall return a pointer to a **struct tm**.

19355 TSF Upon successful completion, *gmtime\_r()* shall return the address of the structure pointed to by the argument *result*. If an error is detected, or UTC is not available, *gmtime\_r()* shall return a null pointer.

## 19358 ERRORS

19359 No errors are defined.

## 19360 EXAMPLES

19361 None.

## 19362 APPLICATION USAGE

19363 The *asctime()*, *ctime()*, *gmtime()*, and *localtime()* functions return values in one of two static objects: a broken-down time structure and an array of **char**. Execution of any of the functions may overwrite the information returned in either of these objects by any of the other functions.

19366 The *gmtime\_r()* function is thread-safe and shall return values in a user-supplied buffer instead of possibly using a static data area that may be overwritten by each call.

## 19368 RATIONALE

19369 None.

## 19370 FUTURE DIRECTIONS

19371 None.

## 19372 SEE ALSO

19373 *asctime()*, *clock()*, *ctime()*, *difftime()*, *localtime()*, *mktime()*, *strftime()*, *strptime()*, *time()*, *utime()*,  
19374 the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

19375 **CHANGE HISTORY**

19376 First released in Issue 1. Derived from Issue 1 of the SVID.

19377 **Issue 4**

19378 In the APPLICATION USAGE section, the list of functions with which this function may interact  
19379 is revised and the wording clarified.

19380 The following change is incorporated for alignment with the ISO C standard:

- 19381 • The type of argument *timer* is changed from **time\_t\*** to **const time\_t\***.

19382 **Issue 5**

19383 A note indicating that the *gmtime()* function need not be reentrant is added to the  
19384 DESCRIPTION.

19385 The *gmtime\_r()* function is included for alignment with the POSIX Threads Extension.

19386 **Issue 6**

19387 The *gmtime\_r()* function is marked as part of the Thread-Safe Functions option.

19388 Extensions beyond the ISO C standard are now marked.

19389 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
19390 its avoidance of possibly using a static data area.

19391 The **restrict** keyword is added to the *gmtime\_r()* prototype for alignment with the  
19392 ISO/IEC 9899:1999 standard.

19393 **NAME**

19394 grantpt — grant access to the slave pseudo-terminal device

19395 **SYNOPSIS**19396 XSI `#include <stdlib.h>`19397 `int grantpt(int fildes);`

19398

19399 **DESCRIPTION**

19400 The *grantpt()* function shall change the mode and ownership of the slave pseudo-terminal  
 19401 device associated with its master pseudo-terminal counterpart. The *fildes* argument is a file  
 19402 descriptor that refers to a master pseudo-terminal device. The user ID of the slave shall be set to  
 19403 the real UID of the calling process and the group ID shall be set to an unspecified group ID. The  
 19404 permission mode of the slave pseudo-terminal shall be set to readable and writable by the  
 19405 owner, and writable by the group.

19406 The behavior of the *grantpt()* function is unspecified if the application has installed a signal  
 19407 handler to catch SIGCHLD signals.

19408 **RETURN VALUE**

19409 Upon successful completion, *grantpt()* shall return 0; otherwise, it shall return -1 and set *errno* to  
 19410 indicate the error.

19411 **ERRORS**19412 The *grantpt()* function may fail if:

|       |          |                                                                                    |  |
|-------|----------|------------------------------------------------------------------------------------|--|
| 19413 | [EBADF]  | The <i>fildes</i> argument is not a valid open file descriptor.                    |  |
| 19414 | [EINVAL] | The <i>fildes</i> argument is not associated with a master pseudo-terminal device. |  |
| 19415 | [EACCES] | The corresponding slave pseudo-terminal device could not be accessed.              |  |

19416 **EXAMPLES**

19417 None.

19418 **APPLICATION USAGE**

19419 None.

19420 **RATIONALE**

19421 None.

19422 **FUTURE DIRECTIONS**

19423 None.

19424 **SEE ALSO**19425 *open()*, *ptsname()*, *unlockpt()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>` |19426 **CHANGE HISTORY**

19427 First released in Issue 4, Version 2.

19428 **Issue 5**

19429 Moved from X/OPEN UNIX extension to BASE.

19430 The last paragraph of the DESCRIPTION is moved from the APPLICATION USAGE section in  
 19431 previous issues.

19432 **NAME**

19433 h\_errno — error return value for network database operations (**LEGACY**)

19434 **SYNOPSIS**

19435 #include <netdb.h>

19436 **DESCRIPTION**

19437 Note that this method of returning errors is used only in connection with legacy functions.

19438 The <**netdb.h**> header provides a declaration of *h\_errno* as a modifiable *l*-value of type **int**.

19439 It is unspecified whether *h\_errno* is a macro or an identifier declared with external linkage. If a  
19440 macro definition is suppressed in order to access an actual object, or a program defines an  
19441 identifier with the name *h\_errno*, the behavior is undefined.

19442 **RETURN VALUE**

19443 None.

19444 **ERRORS**

19445 No errors are defined.

19446 **EXAMPLES**

19447 None.

19448 **APPLICATION USAGE**

19449 Applications should obtain the definition of *h\_errno* by the inclusion of the <**netdb.h**> header.

19450 The practice of defining *h\_errno* in a program as an **extern int h\_errno** is obsolescent.

19451 **RATIONALE**

19452 None.

19453 **FUTURE DIRECTIONS**

19454 *h\_errno* may be withdrawn in a future version.

19455 **SEE ALSO**

19456 *endhostent()*, *errno*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**netdb.h**>

19457 **CHANGE HISTORY**

19458 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19459 **NAME**

19460 hcreate, hdestroy, hsearch — manage hash search table

19461 **SYNOPSIS**

```

19462 xSI #include <search.h>
19463 int hcreate(size_t nel);
19464 void hdestroy(void);
19465 ENTRY *hsearch(ENTRY item, ACTION action);
19466

```

19467 **DESCRIPTION**19468 The *hcreate()*, *hdestroy()*, and *hsearch()* functions manage hash search tables.

19469 The *hcreate()* function allocates sufficient space for the table, and the application shall ensure it is called before *hsearch()* is used. The *nel* argument is an estimate of the maximum number of entries that the table shall contain. This number may be adjusted upward by the algorithm in order to obtain certain mathematically favorable circumstances.

19473 The *hdestroy()* function disposes of the search table, and may be followed by another call to *hcreate()*. After the call to *hdestroy()*, the data can no longer be considered accessible.

19475 The *hsearch()* function is a hash-table search routine. It shall return a pointer into a hash table indicating the location at which an entry can be found. The *item* argument is a structure of type **ENTRY** (defined in the *<search.h>* header) containing two pointers: *item.key* points to the comparison key (a **char\***), and *item.data* (a **void\***) points to any other data to be associated with that key. The comparison function used by *hsearch()* is *strcmp()*. The *action* argument is a member of an enumeration type **ACTION** indicating the disposition of the entry if it cannot be found in the table. **ENTER** indicates that the item should be inserted in the table at an appropriate point. **FIND** indicates that no entry should be made. Unsuccessful resolution is indicated by the return of a null pointer.

19484 These functions need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

19486 **RETURN VALUE**

19487 The *hcreate()* function shall return 0 if it cannot allocate sufficient space for the table; otherwise, it shall return non-zero.

19489 The *hdestroy()* function shall return no value.

19490 The *hsearch()* function shall return a null pointer if either the action is **FIND** and the item could not be found or the action is **ENTER** and the table is full.

19492 **ERRORS**

19493 The *hcreate()* and *hsearch()* functions may fail if:

19494 [ENOMEM] Insufficient storage space is available.

19495 **EXAMPLES**

19496 The following example reads in strings followed by two numbers and stores them in a hash table, discarding duplicates. It then reads in strings and finds the matching entry in the hash table and prints it out.

```

19499 #include <stdio.h>
19500 #include <search.h>
19501 #include <string.h>
19502 struct info { /* This is the info stored in the table */
19503 int age, room; /* other than the key. */

```

```

19504 };
19505 #define NUM_EMPL 5000 /* # of elements in search table. */
19506 int main(void)
19507 {
19508 char string_space[NUM_EMPL*20]; /* Space to store strings. */
19509 struct info info_space[NUM_EMPL]; /* Space to store employee info. */
19510 char *str_ptr = string_space; /* Next space in string_space. */
19511 struct info *info_ptr = info_space;
19512 /* Next space in info_space. */
19513 ENTRY item;
19514 ENTRY *found_item; /* Name to look for in table. */
19515 char name_to_find[30];
19516
19516 int i = 0;
19517
19517 /* Create table; no error checking is performed. */
19518 (void) hcreate(NUM_EMPL);
19519 while (scanf("%s%d%d", str_ptr, &info_ptr->age,
19520 &info_ptr->room) != EOF && i++ < NUM_EMPL) {
19521
19521 /* Put information in structure, and structure in item. */
19522 item.key = str_ptr;
19523 item.data = info_ptr;
19524 str_ptr += strlen(str_ptr) + 1;
19525 info_ptr++;
19526
19526 /* Put item into table. */
19527 (void) hsearch(item, ENTER);
19528 }
19529
19529 /* Access table. */
19530 item.key = name_to_find;
19531 while (scanf("%s", item.key) != EOF) {
19532 if ((found_item = hsearch(item, FIND)) != NULL) {
19533
19533 /* If item is in the table. */
19534 (void)printf("found %s, age = %d, room = %d\n",
19535 found_item->key,
19536 ((struct info *)found_item->data)->age,
19537 ((struct info *)found_item->data)->room);
19538 } else
19539 (void)printf("no such employee %s\n", name_to_find);
19540 }
19541 return 0;
19542 }

```

#### 19543 APPLICATION USAGE

19544 The *hcreate()* and *hsearch()* functions may use *malloc()* to allocate space.

#### 19545 RATIONALE

19546 None.

19547 **FUTURE DIRECTIONS**

19548           None.

19549 **SEE ALSO**

19550           *bsearch()*, *lsearch()*, *malloc()*, *strcmp()*, *tsearch()*, the Base Definitions volume of  
19551           IEEE Std. 1003.1-200x, <**search.h**>

19552 **CHANGE HISTORY**

19553           First released in Issue 1. Derived from Issue 1 of the SVID.

19554 **Issue 4**

19555           In the SYNOPSIS section, the type of argument *nel* in the declaration of *hcreate()* is changed from  
19556           **unsigned** to **size\_t**, and the argument list is explicitly defined as **void** in the declaration of  
19557           *hdestroy()*.

19558           In the DESCRIPTION, the type of the comparison key is explicitly defined as **char\***, the type of  
19559           *item.data* is explicitly defined as **void\***, and a statement is added indicating that *hsearch()* uses  
19560           *strcmp()* as the comparison function.

19561           In the EXAMPLES section, the sample code is updated to use ISO C standard syntax.

19562           An ERRORS section is added and [ENOMEM] is defined as an error that may be returned by  
19563           *hsearch()* and *hcreate()*.

19564 **Issue 6**

19565           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

19566 **NAME**

19567 htonl, htons, ntohl, ntohs — convert values between host and network byte order

19568 **SYNOPSIS**

19569 #include <arpa/inet.h>

19570 uint32\_t htonl(uint32\_t *hostlong*);

19571 uint16\_t htons(uint16\_t *hostshort*);

19572 uint32\_t ntohl(uint32\_t *netlong*);

19573 uint16\_t ntohs(uint16\_t *netshort*);

19574 **DESCRIPTION**

19575 These functions shall convert 16-bit and 32-bit quantities between network byte order and host  
19576 byte order.

19577 On some implementations, these functions are defined as macros.

19578 The **uint32\_t** and **uint16\_t** types shall be defined as described in <inttypes.h>.

19579 **RETURN VALUE**

19580 The *htonl()* and *htons()* functions shall return the argument value converted from host to  
19581 network byte order.

19582 The *ntohl()* and *ntohs()* functions shall return the argument value converted from network to  
19583 host byte order.

19584 **ERRORS**

19585 No errors are defined.

19586 **EXAMPLES**

19587 None.

19588 **APPLICATION USAGE**

19589 These functions are most often used in conjunction with IPv4 addresses and ports as returned by  
19590 *gethostent()* and *getservent()*.

19591 **RATIONALE**

19592 None.

19593 **FUTURE DIRECTIONS**

19594 None.

19595 **SEE ALSO**

19596 *endhostent()*, *endservent()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <inttypes.h>,  
19597 <arpa/inet.h>

19598 **CHANGE HISTORY**

19599 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19600 **NAME**

19601       htons — convert values between host and network byte order

19602 **SYNOPSIS**

19603       #include <arpa/inet.h>

19604       uint16\_t htons(uint16\_t *hostshort*);

19605 **DESCRIPTION**

19606       Refer to *htonl()*.

19607 **NAME**

19608 hypot, hypotf, hypotl — Euclidean distance function

19609 **SYNOPSIS**

```
19610 xSI #include <math.h>
19611 double hypot(double x, double y);
19612 float hypotf(float x, float y);
19613 long double hypotl(long double x, long double y);
19614
```

19615 **DESCRIPTION**19616 These functions shall compute the value of the square root of  $x^2+y^2$ .

19617 An application wishing to check for error situations should set *errno* to 0 before calling *hypot()*.  
 19618 If *errno* is non-zero on return, or the return value is HUGE\_VAL or NaN, an error has occurred.

19619 **RETURN VALUE**19620 Upon successful completion, these functions shall return the length of the hypotenuse of a  
19621 right-angled triangle with sides of length *x* and *y*.19622 If the result would cause overflow, HUGE\_VAL shall be returned and *errno* may be set to  
19623 [ERANGE].19624 If *x* or *y* is NaN, NaN shall be returned. and *errno* may be set to [EDOM].19625 If the correct result would cause underflow, 0 shall be returned and *errno* may be set to  
19626 [ERANGE].19627 **ERRORS**

19628 These functions may fail if:

19629 [EDOM] The value of *x* or *y* is NaN.

19630 [ERANGE] The result overflows or underflows.

19631 No other errors shall occur.

19632 **EXAMPLES**

19633 None.

19634 **APPLICATION USAGE**

19635 The *hypot()* function takes precautions against overflow during intermediate steps of the  
 19636 computation. If the calculated result would still overflow a double, then *hypot()* shall return  
 19637 HUGE\_VAL.

19638 **RATIONALE**

19639 None.

19640 **FUTURE DIRECTIONS**

19641 None.

19642 **SEE ALSO**19643 *isnan()*, *sqrt()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>19644 **CHANGE HISTORY**

19645 First released in Issue 1. Derived from Issue 1 of the SVID.

19646 **Issue 4**

19647       References to *matherr()* are removed.

19648       The RETURN VALUE and ERRORS sections are substantially rewritten to rationalize error  
19649       handling in the mathematics functions.

19650 **Issue 5**

19651       The DESCRIPTION is updated to indicate how an application should check for an error. This  
19652       text was previously published in the APPLICATION USAGE section.

19653 **Issue 6**

19654       The *hypotf()* and *hypotl()* functions are added for alignment with the ISO/IEC 9899:1999  
19655       standard.

19656 **NAME**

19657 iconv — codeset conversion function

19658 **SYNOPSIS**19659 XSI 

```
#include <iconv.h>
```

```
19660 size_t iconv(iconv_t cd, char **restrict inbuf,
19661 size_t *restrict inbytesleft, char **restrict outbuf,
19662 size_t *restrict outbytesleft);
19663
```

19664 **DESCRIPTION**

19665 The *iconv()* function shall convert the sequence of characters from one codeset, in the array  
 19666 specified by *inbuf*, into a sequence of corresponding characters in another codeset, in the array  
 19667 specified by *outbuf*. The codesets are those specified in the *iconv\_open()* call that returned the  
 19668 conversion descriptor, *cd*. The *inbuf* argument points to a variable that points to the first  
 19669 character in the input buffer and *inbytesleft* indicates the number of bytes to the end of the buffer  
 19670 to be converted. The *outbuf* argument points to a variable that points to the first available byte in  
 19671 the output buffer and *outbytesleft* indicates the number of the available bytes to the end of the  
 19672 buffer.

19673 For state-dependent encodings, the conversion descriptor *cd* is placed into its initial shift state by  
 19674 a call for which *inbuf* is a null pointer, or for which *inbuf* points to a null pointer. When *iconv()*  
 19675 is called in this way, and if *outbuf* is not a null pointer or a pointer to a null pointer, and *outbytesleft*  
 19676 points to a positive value, *iconv()* shall place, into the output buffer, the byte sequence to change  
 19677 the output buffer to its initial shift state. If the output buffer is not large enough to hold the  
 19678 entire reset sequence, *iconv()* shall fail and set *errno* to [E2BIG]. Subsequent calls with *inbuf* as  
 19679 other than a null pointer or a pointer to a null pointer cause the conversion to take place from  
 19680 the current state of the conversion descriptor.

19681 If a sequence of input bytes does not form a valid character in the specified codeset, conversion  
 19682 stops after the previous successfully converted character. If the input buffer ends with an  
 19683 incomplete character or shift sequence, conversion stops after the previous successfully  
 19684 converted bytes. If the output buffer is not large enough to hold the entire converted input,  
 19685 conversion stops just prior to the input bytes that would cause the output buffer to overflow.  
 19686 The variable pointed to by *inbuf* is updated to point to the byte following the last byte  
 19687 successfully used in the conversion. The value pointed to by *inbytesleft* is decremented to reflect  
 19688 the number of bytes still not converted in the input buffer. The variable pointed to by *outbuf* is  
 19689 updated to point to the byte following the last byte of converted output data. The value pointed  
 19690 to by *outbytesleft* is decremented to reflect the number of bytes still available in the output buffer.  
 19691 For state-dependent encodings, the conversion descriptor is updated to reflect the shift state in  
 19692 effect at the end of the last successfully converted byte sequence.

19693 If *iconv()* encounters a character in the input buffer that is valid, but for which an identical  
 19694 character does not exist in the target codeset, *iconv()* performs an implementation-defined  
 19695 conversion on this character.

19696 **RETURN VALUE**

19697 The *iconv()* function shall update the variables pointed to by the arguments to reflect the extent  
 19698 of the conversion and return the number of non-identical conversions performed. If the entire  
 19699 string in the input buffer is converted, the value pointed to by *inbytesleft* shall be 0. If the input  
 19700 conversion is stopped due to any conditions mentioned above, the value pointed to by *inbytesleft*  
 19701 shall be non-zero and *errno* shall be set to indicate the condition. If an error occurs *iconv()* shall  
 19702 return (*size\_t*)-1 and set *errno* to indicate the error.

19703 **ERRORS**

19704 The *iconv()* function shall fail if:

19705 [EILSEQ] Input conversion stopped due to an input byte that does not belong to the  
19706 input codeset.

19707 [E2BIG] Input conversion stopped due to lack of space in the output buffer.

19708 [EINVAL] Input conversion stopped due to an incomplete character or shift sequence at  
19709 the end of the input buffer.

19710 The *iconv()* function may fail if:

19711 [EBADF] The *cd* argument is not a valid open conversion descriptor.

19712 **EXAMPLES**

19713 None.

19714 **APPLICATION USAGE**

19715 The *inbuf* argument indirectly points to the memory area which contains the conversion input  
19716 data. The *outbuf* argument indirectly points to the memory area which is to contain the result of  
19717 the conversion. The objects indirectly pointed to by *inbuf* and *outbuf* are not restricted to  
19718 containing data that is directly representable in the ISO C standard language **char** data type. The  
19719 type of *inbuf* and *outbuf*, **char\*\***, does not imply that the objects pointed to are interpreted as  
19720 null-terminated C strings or arrays of characters. Any interpretation of a byte sequence that  
19721 represents a character in a given character set encoding scheme is done internally within the  
19722 codeset converters. For example, the area pointed to indirectly by *inbuf* and/or *outbuf* can  
19723 contain all zero octets that are not interpreted as string terminators but as coded character data  
19724 according to the respective codeset encoding scheme. The type of the data (**char**, **short**, **long**, and  
19725 so on) read or stored in the objects is not specified, but may be inferred for both the input and  
19726 output data by the converters determined by the *fromcode* and *toctype* arguments of *iconv\_open()*.

19727 Regardless of the data type inferred by the converter, the size of the remaining space in both  
19728 input and output objects (the *inbytesleft* and *outbytesleft* arguments) is always measured in bytes.

19729 For implementations that support the conversion of state-dependent encodings, the conversion  
19730 descriptor must be able to accurately reflect the shift-state in effect at the end of the last  
19731 successful conversion. It is not required that the conversion descriptor itself be updated, which  
19732 would require it to be a pointer type. Thus, implementations are free to implement the  
19733 descriptor as a handle (other than a pointer type) by which the conversion information can be  
19734 accessed and updated.

19735 **RATIONALE**

19736 None.

19737 **FUTURE DIRECTIONS**

19738 None.

19739 **SEE ALSO**

19740 *iconv\_open()*, *iconv\_close()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**iconv.h**>

19741 **CHANGE HISTORY**

19742 First released in Issue 4. Derived from the HP-UX Manual.

19743 **Issue 6**

19744 The SYNOPSIS has been corrected to align with the <**iconv.h**> reference page.

19745 The **restrict** keyword is added to the *iconv()* prototype for alignment with the  
19746 ISO/IEC 9899:1999 standard.

19747 **NAME**

19748 iconv\_close — codeset conversion deallocation function

19749 **SYNOPSIS**

19750 XSI #include &lt;iconv.h&gt;

19751 int iconv\_close(iconv\_t *cd*);

19752

19753 **DESCRIPTION**19754 The *iconv\_close()* function deallocates the conversion descriptor *cd* and all other associated  
19755 resources allocated by *iconv\_open()*.19756 If a file descriptor is used to implement the type **iconv\_t**, that file descriptor is closed.19757 **RETURN VALUE**19758 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to  
19759 indicate the error.19760 **ERRORS**19761 The *iconv\_close()* function may fail if:

19762 [EBADF] The conversion descriptor is invalid.

19763 **EXAMPLES**

19764 None.

19765 **APPLICATION USAGE**

19766 None.

19767 **RATIONALE**

19768 None.

19769 **FUTURE DIRECTIONS**

19770 None.

19771 **SEE ALSO**19772 *iconv()*, *iconv\_open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**iconv.h**>19773 **CHANGE HISTORY**

19774 First released in Issue 4. Derived from the HP-UX Manual.

19775 **NAME**

19776 iconv\_open — codeset conversion allocation function

19777 **SYNOPSIS**19778 XSI `#include <iconv.h>`19779 `iconv_t iconv_open(const char *tocode, const char *fromcode);`

19780

19781 **DESCRIPTION**

19782 The *iconv\_open()* function shall return a conversion descriptor that describes a conversion from  
 19783 the codeset specified by the string pointed to by the *fromcode* argument to the codeset specified  
 19784 by the string pointed to by the *tocode* argument. For state-dependent encodings, the conversion  
 19785 descriptor is in a codeset-dependent initial shift state, ready for immediate use with *iconv()*.

19786 Settings of *fromcode* and *tocode* and their permitted combinations are implementation-defined.19787 A conversion descriptor remains valid until it is closed by *iconv\_close()* or an implicit close.

19788 If a file descriptor is used to implement conversion descriptors, the FD\_CLOEXEC flag shall be  
 19789 set; see <fcntl.h>.

19790 **RETURN VALUE**

19791 Upon successful completion, *iconv\_open()* shall return a conversion descriptor for use on  
 19792 subsequent calls to *iconv()*. Otherwise, *iconv\_open()* shall return (**iconv\_t**)-1 and set *errno* to  
 19793 indicate the error.

19794 **ERRORS**19795 The *iconv\_open()* function may fail if:

19796 [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

19797 [ENFILE] Too many files are currently open in the system.

19798 [ENOMEM] Insufficient storage space is available.

19799 [EINVAL] The conversion specified by *fromcode* and *tocode* is not supported by the  
 19800 implementation.

19801 **EXAMPLES**

19802 None.

19803 **APPLICATION USAGE**

19804 Some implementations of *iconv\_open()* use *malloc()* to allocate space for internal buffer areas.  
 19805 The *iconv\_open()* function may fail if there is insufficient storage space to accommodate these  
 19806 buffers.

19807 Portable applications must assume that conversion descriptors are not valid after a call to one of  
 19808 the *exec* functions.

19809 **RATIONALE**

19810 None.

19811 **FUTURE DIRECTIONS**

19812 None.

19813 **SEE ALSO**19814 *iconv()*, *iconv\_close()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <fcntl.h>, <iconv.h>

19815 **CHANGE HISTORY**

19816 First released in Issue 4. Derived from the HP-UX Manual.

19817 **NAME**

19818 if\_freenameindex — free memory allocated by *ifnameindex()*

19819 **SYNOPSIS**

19820 #include <net/if.h>

19821 void if\_freenameindex(struct if\_nameindex \*ptr);

19822 **DESCRIPTION**

19823 The *if\_freenameindex()* function shall free the memory allocated by *if\_nameindex*. The *ptr* |  
19824 argument shall be a pointer that was returned by *if\_nameindex*. After *if\_freenameindex()* has been |  
19825 called, the application should not use the array of which *ptr* is the address.

19826 **RETURN VALUE**

19827 None.

19828 **ERRORS**

19829 No errors are defined.

19830 **EXAMPLES**

19831 None.

19832 **APPLICATION USAGE**

19833 None.

19834 **RATIONALE**

19835 None.

19836 **FUTURE DIRECTIONS**

19837 None.

19838 **SEE ALSO**

19839 *getsockopt()*, *if\_indextoname()*, *if\_nameindex()*, *if\_nametoindex()*, *setsockopt()*, the Base Definitions |  
19840 volume of IEEE Std. 1003.1-200x, <net/if.h>

19841 **CHANGE HISTORY**

19842 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19843 **NAME**

19844 if\_indextoname — map a network interface index to its corresponding name

19845 **SYNOPSIS**

19846 #include <net/if.h>

19847 char \*if\_indextoname(unsigned *ifindex*, char \**ifname*);

19848 **DESCRIPTION**

19849 The *if\_indextoname()* function shall map an interface index to its corresponding name.

19850 When this function is called, *ifname* shall point to a buffer of at least {IFNAMSIZ} bytes. The  
19851 function shall place in this buffer the name of the interface with index *ifindex*.

19852 **RETURN VALUE**

19853 If *ifindex* is an interface index, then the function shall return the value supplied in *ifname*, which  
19854 points to a buffer now containing the interface name. Otherwise, the function shall return a  
19855 NULL pointer and set *errno* to indicate the error.

19856 **ERRORS**

19857 The *if\_indextoname()* function shall fail if:

19858 [ENXIO] The interface does not exist.

19859 **EXAMPLES**

19860 None.

19861 **APPLICATION USAGE**

19862 None.

19863 **RATIONALE**

19864 None.

19865 **FUTURE DIRECTIONS**

19866 None.

19867 **SEE ALSO**

19868 *getsockopt()*, *if\_freenameindex()*, *if\_nameindex()*, *if\_nametoindex()*, *setsockopt()*, the Base  
19869 Definitions volume of IEEE Std. 1003.1-200x, <net/if.h>

19870 **CHANGE HISTORY**

19871 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19872 **NAME**

19873 if\_nameindex — return all network interface names and indexes

19874 **SYNOPSIS**

19875 #include <net/if.h>

19876 struct if\_nameindex \*if\_nameindex(void);

19877 **DESCRIPTION**

19878 The *if\_nameindex()* function shall return an array of *if\_nameindex* structures, one structure per  
19879 interface. The end of the array is indicated by a structure with an *if\_index* field of zero and an  
19880 *if\_name* field of NULL.

19881 Applications should call *if\_freenameindex()* to release the memory that may be dynamically  
19882 allocated by this function, after they have finished using it.

19883 **RETURN VALUE**

19884 Array of structures identifying local interfaces. A NULL pointer is returned upon an error, with  
19885 *errno* set to indicate the error.

19886 **ERRORS**

19887 The *if\_nameindex()* function may fail if:

19888 [ENOBUFS] Insufficient resources are available to complete the function.

19889 **EXAMPLES**

19890 None.

19891 **APPLICATION USAGE**

19892 None.

19893 **RATIONALE**

19894 None.

19895 **FUTURE DIRECTIONS**

19896 None.

19897 **SEE ALSO**

19898 *getsockopt()*, *if\_freenameindex()*, *if\_indextoname()*, *if\_nametoindex()*, *setsockopt()*, the Base  
19899 Definitions volume of IEEE Std. 1003.1-200x, <net/if.h>

19900 **CHANGE HISTORY**

19901 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

19902 **NAME**

19903 if\_nametoindex — map a network interface name to its corresponding index |

19904 **SYNOPSIS**

19905 #include <net/if.h>

19906 unsigned if\_nametoindex(const char \*ifname); |

19907 **DESCRIPTION**

19908 The *if\_nametoindex()* function shall return the interface index corresponding to name *ifname*. |

19909 **RETURN VALUE**

19910 The corresponding index if *ifname* is the name of an interface; otherwise, zero. |

19911 **ERRORS**

19912 No errors are defined. |

19913 **EXAMPLES**

19914 None.

19915 **APPLICATION USAGE**

19916 None.

19917 **RATIONALE**

19918 None.

19919 **FUTURE DIRECTIONS**

19920 None.

19921 **SEE ALSO**

19922 *getsockopt()*, *if\_freenameindex()*, *if\_indextoname()*, *if\_nameindex()*, *setsockopt()*, the Base |

19923 Definitions volume of IEEE Std. 1003.1-200x, <net/if.h> |

19924 **CHANGE HISTORY**

19925 First released in Issue 6. Derived from the XNS, Issue 5.2 specification. |

19926 **NAME**

19927       ilogb, ilogbf, ilogbl — return an unbiased exponent

19928 **SYNOPSIS**

19929       #include <math.h>

19930       int ilogb(double x);

19931       int ilogbf(float x);

19932       int ilogbl(long double x);

19933 **DESCRIPTION**

19934       These functions shall return the exponent part of  $x$ . Formally, the return value is the integral part of  $\log_r |x|$  as a signed integral value, for non-zero  $x$ , where  $r$  is the radix of the machine's floating point arithmetic, which is the value of FLT\_RADIX defined in <float.h>.

19937 **RETURN VALUE**

19938       Upon successful completion, these functions shall return the exponent part of  $x$  as a signed integer value.

19940       If  $x$  is zero, these functions shall return the value FP\_ILOGB0; if  $x$  is infinite, they shall return the value {INT\_MAX}; if  $x$  is a NaN, they shall return the value FP\_ILOGBNAN; otherwise, they shall be equivalent to calling the corresponding *logb()* function and casting the returned value to type **int**.

19944 **ERRORS**

19945       These functions may fail if:

19946       [ERANGE]       The value of  $x$  is 0.

19947 **EXAMPLES**

19948       None.

19949 **APPLICATION USAGE**

19950       None.

19951 **RATIONALE**

19952       None.

19953 **FUTURE DIRECTIONS**

19954       None.

19955 **SEE ALSO**

19956       *logb()*, *scalb()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <float.h>, <math.h>

19957 **CHANGE HISTORY**

19958       First released in Issue 4, Version 2.

19959 **Issue 5**

19960       Moved from X/OPEN UNIX extension to BASE.

19961 **Issue 6**

19962       The *ilogbf()* and *ilogbl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

19964 **NAME**

19965           imaxabs — return absolute value

19966 **SYNOPSIS**

19967           #include &lt;inttypes.h&gt;

19968           intmax\_t imaxabs(intmax\_t j);

19969 **DESCRIPTION**

19970 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
19971       conflict between the requirements described here and the ISO C standard is unintentional. This  
19972       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

19973       The *imaxabs()* function shall compute the absolute value of an integer *j*. If the result cannot be  
19974       represented, the behavior is undefined.

19975 **RETURN VALUE**19976       The *imaxabs()* function shall return the absolute value.19977 **ERRORS**

19978       No errors are defined.

19979 **EXAMPLES**

19980       None.

19981 **APPLICATION USAGE**

19982       The absolute value of the most negative number cannot be represented in two's complement.

19983 **RATIONALE**

19984       None.

19985 **FUTURE DIRECTIONS**

19986       None.

19987 **SEE ALSO**19988       *imaxdiv()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <inttypes.h>19989 **CHANGE HISTORY**

19990       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

19991 **NAME**

19992 imaxdiv — return quotient and remainder

19993 **SYNOPSIS**

19994 #include &lt;inttypes.h&gt;

19995 imaxdiv\_t imaxdiv(intmax\_t numer, intmax\_t denom);

19996 **DESCRIPTION**

19997 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
19998 conflict between the requirements described here and the ISO C standard is unintentional. This  
19999 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

20000 The *imaxdiv()* function shall compute *numer / denom* and *numer % denom* in a single operation.20001 **RETURN VALUE**

20002 The *imaxdiv()* function shall return a structure of type **imaxdiv\_t**, comprising both the quotient  
20003 and the remainder. The structure shall contain (in either order) the members *quot* (the quotient)  
20004 and *rem* (the remainder), each of which has type **intmax\_t**.

20005 If either part of the result cannot be represented, the behavior is undefined.

20006 **ERRORS**

20007 No errors are defined.

20008 **EXAMPLES**

20009 None.

20010 **APPLICATION USAGE**

20011 None.

20012 **RATIONALE**

20013 None.

20014 **FUTURE DIRECTIONS**

20015 None.

20016 **SEE ALSO**20017 *imaxabs()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**inttypes.h**>20018 **CHANGE HISTORY**

20019 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

20020 **NAME**20021 index — character string operations (**LEGACY**)20022 **SYNOPSIS**20023 XSI `#include <strings.h>`20024 `char *index(const char *s, int c);`

20025

20026 **DESCRIPTION**20027 The *index()* function is identical to *strchr()*.20028 **RETURN VALUE**20029 See *strchr()*.20030 **ERRORS**20031 See *strchr()*.20032 **EXAMPLES**

20033 None.

20034 **APPLICATION USAGE**20035 *strchr()* is preferred over this function.20036 For maximum portability, it is recommended to replace the function call to *index()* as follows:20037 `#define index(a,b) strchr((a),(b))`20038 **RATIONALE**

20039 None.

20040 **FUTURE DIRECTIONS**

20041 This function may be withdrawn in a future version.

20042 **SEE ALSO**20043 *strchr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<strings.h>`20044 **CHANGE HISTORY**

20045 First released in Issue 4, Version 2.

20046 **Issue 5**

20047 Moved from X/OPEN UNIX extension to BASE.

20048 **Issue 6**

20049 This function is marked LEGACY.

## 20050 NAME

20051 inet\_addr, inet\_lnaof (LEGACY), inet\_makeaddr (LEGACY), inet\_netof (LEGACY),  
 20052 inet\_network (LEGACY), inet\_ntoa — IPv4 address manipulation

## 20053 SYNOPSIS

```
20054 #include <arpa/inet.h>

20055 in_addr_t inet_addr(const char *cp);
20056 in_addr_t inet_lnaof(struct in_addr in);
20057 struct in_addr inet_makeaddr(in_addr_t net, in_addr_t lna);
20058 in_addr_t inet_netof(struct in_addr in);
20059 in_addr_t inet_network(const char *cp);
20060 char *inet_ntoa(struct in_addr in);
```

## 20061 DESCRIPTION

20062 The *inet\_addr()* function shall convert the string pointed to by *cp*, in the standard IPv4 dotted  
 20063 decimal notation, to an integer value suitable for use as an Internet address.

20064 The *inet\_lnaof()* function shall take an Internet host address specified by *in* and extract the local  
 20065 network address part, in host byte order.

20066 The *inet\_makeaddr()* function shall take the Internet network number specified by *net* and the  
 20067 local network address specified by *lna*, both in host byte order, and construct an Internet address  
 20068 from them.

20069 The *inet\_netof()* function shall take an Internet host address specified by *in* and extract the  
 20070 network number part, in host byte order.

20071 The *inet\_network()* function shall convert the string pointed to by *cp*, in the standard IPv4 dotted  
 20072 decimal notation, to an integer value suitable for use as an Internet network number.

20073 The *inet\_ntoa()* function shall convert the Internet host address specified by *in* to a string in the  
 20074 Internet standard dot notation.

20075 All Internet addresses shall be returned in network order (bytes ordered from left to right).

20076 Values specified using IPv4 dotted decimal notation take one of the following forms:

20077 a.b.c.d    When four parts are specified, each is interpreted as a byte of data and assigned,  
 20078 from left to right, to the four bytes of an Internet address.

20079 a.b.c      When a three-part address is specified, the last part is interpreted as a 16-bit  
 20080 quantity and placed in the rightmost two bytes of the network address. This makes  
 20081 the three-part address format convenient for specifying Class B network addresses  
 20082 as **128.net.host**.

20083 a.b        When a two-part address is supplied, the last part is interpreted as a 24-bit  
 20084 quantity and placed in the rightmost three bytes of the network address. This  
 20085 makes the two-part address format convenient for specifying Class A network  
 20086 addresses as **net.host**.

20087 a          When only one part is given, the value is stored directly in the network address  
 20088 without any byte rearrangement.

20089 All numbers supplied as parts in IPv4 dotted decimal notation may be decimal, octal, or  
 20090 hexadecimal, as specified in the ISO C standard (that is, a leading "0x" or "0X" implies  
 20091 hexadecimal; otherwise, a leading '0' implies octal; otherwise, the number is interpreted as  
 20092 decimal).

**20093 RETURN VALUE**

20094 Upon successful completion, *inet\_addr()* shall return the Internet address. Otherwise, it shall  
20095 return (**in\_addr\_t**)(-1).

20096 The *inet\_lnaof()* function shall return the local network address part.

20097 The *inet\_makeaddr()* function shall return the constructed Internet address.

20098 The *inet\_netof()* function shall return the network number.

20099 Upon successful completion, *inet\_network()* shall return the converted Internet network number.  
20100 Otherwise, it shall return (**in\_addr\_t**)(-1).

20101 The *inet\_ntoa()* function shall return a pointer to the network address in Internet standard dot  
20102 notation.

**20103 ERRORS**

20104 No errors are defined.

**20105 EXAMPLES**

20106 None.

**20107 APPLICATION USAGE**

20108 The *inet\_lnaof()*, *inet\_makeaddr()*, *inet\_netof()*, and *inet\_network()* functions are marked LEGACY  
20109 and should not be used by new applications.

20110 The return value of *inet\_ntoa()* may point to static data that may be overwritten by subsequent  
20111 calls to *inet\_ntoa()*.

**20112 RATIONALE**

20113 None.

**20114 FUTURE DIRECTIONS**

20115 The *inet\_lnaof()*, *inet\_makeaddr()*, *inet\_netof()*, and *inet\_network()* functions may be withdrawn in  
20116 a future version.

**20117 SEE ALSO**

20118 *endhostent()*, *endnetent()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <arpa/inet.h>

**20119 CHANGE HISTORY**

20120 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

20121 **NAME**

20122           inet\_lnaof — IPv4 address manipulation

20123 **SYNOPSIS**

20124           #include <arpa/inet.h>

20125           in\_addr\_t inet\_lnaof(struct in\_addr *in*);

20126 **DESCRIPTION**

20127           Refer to *inet\_addr()*.

20128 **NAME**

20129       inet\_makeaddr — IPv4 address manipulation

20130 **SYNOPSIS**

20131       #include <arpa/inet.h>

20132       struct in\_addr inet\_makeaddr(in\_addr\_t net, in\_addr\_t lna);

20133 **DESCRIPTION**

20134       Refer to *inet\_addr()*.

20135 **NAME**

20136       inet\_netof — IPv4 address manipulation

20137 **SYNOPSIS**

20138       #include <arpa/inet.h>

20139       in\_addr\_t inet\_netof(struct in\_addr *in*);

20140 **DESCRIPTION**

20141       Refer to *inet\_addr()*.

20142 **NAME**

20143       inet\_network — IPv4 address manipulation

20144 **SYNOPSIS**

20145       #include <arpa/inet.h>

20146       in\_addr\_t inet\_network(const char \*cp);

20147 **DESCRIPTION**

20148       Refer to *inet\_addr()*.

20149 **NAME**

20150       inet\_ntoa — IPv4 address manipulation

20151 **SYNOPSIS**

20152       #include <arpa/inet.h>

20153       char \*inet\_ntoa(struct in\_addr *in*);

20154 **DESCRIPTION**

20155       Refer to *inet\_addr()*.

20156 **NAME**

20157 inet\_ntop, inet\_pton — convert IPv4 and IPv6 addresses between binary and text form

20158 **SYNOPSIS**

20159 IP6 `#include <arpa/inet.h>`

```
20160 const char *inet_ntop(int af, const void *restrict src,
20161 char *restrict dst, socklen_t size);
20162 int inet_pton(int af, const char *restrict src, void *restrict dst);
20163
```

20164 **Notes to Reviewers**

20165 *This section with side shading will not appear in the final copy. - Ed.*

20166 D3, XSH, ERN 330 (AI 2000-05-024) To supply editing instructions regarding making these  
20167 functions mandatory and shading IPv6-specific information.

20168 **DESCRIPTION**

20169 The *inet\_ntop()* function converts a numeric address into a text string suitable for presentation.  
20170 The *af* argument specifies the family of the address. This can be AF\_INET or AF\_INET6. The *src*  
20171 argument points to a buffer holding an IPv4 address if the *af* argument is AF\_INET, or an IPv6  
20172 address if the *af* argument is AF\_INET6. The *dst* argument points to a buffer where the function  
20173 stores the resulting text string; it shall not be NULL. The *size* argument specifies the size of this  
20174 buffer, which shall be large enough to hold the text string (INET\_ADDRSTRLEN characters for  
20175 IPv4, INET6\_ADDRSTRLEN characters for IPv6).

20176 The *inet\_pton()* function converts an address in its standard text presentation form into its  
20177 numeric binary form. The *af* argument specifies the family of the address. The AF\_INET and  
20178 AF\_INET6 address families are supported. The *src* argument points to the string being passed in.  
20179 The *dst* argument points to a buffer into which the function stores the numeric address; this shall  
20180 be large enough to hold the numeric address (32 bits for AF\_INET, 128 bits for AF\_INET6).

20181 If the *af* argument of *inet\_pton()* is AF\_INET, the *src* string shall be in the standard IPv4 dotted-  
20182 decimal form:

20183 `ddd . ddd . ddd . ddd`

20184 where "ddd" is a one to three digit decimal number between 0 and 255 (see *inet\_addr()*). The  
20185 *inet\_pton()* function does not accept other formats (such as the octal numbers, hexadecimal  
20186 numbers, and fewer than four numbers that *inet\_addr()* accepts).

20187 If the *af* argument of *inet\_pton()* is AF\_INET6, the *src* string shall be in one of the following  
20188 standard IPv6 text forms:

- 20189 1. The preferred form is "`x:x:x:x:x:x:x:x`", where the 'x's are the hexadecimal values  
20190 of the eight 16-bit pieces of the address. Leading zeros in individual fields can be omitted,  
20191 but there shall be at least one numeral in every field.
- 20192 2. A string of contiguous zero fields in the preferred form can be shown as "`:::`". The "`:::`"  
20193 can only appear once in an address. Unspecified addresses ("`0:0:0:0:0:0:0:0`") may  
20194 be represented simply as "`:::`".
- 20195 3. A third form that is sometimes more convenient when dealing with a mixed environment  
20196 of IPv4 and IPv6 nodes is "`x:x:x:x:x:x:d.d.d.d`", where the 'x's are the  
20197 hexadecimal values of the six high-order 16-bit pieces of the address, and the 'd's are the  
20198 decimal values of the four low-order 8-bit pieces of the address (standard IPv4  
20199 representation).

20200           **Note:**       A more extensive description of the standard representations of IPv6 addresses can  
20201                           be found in RFC 2373.

20202 **RETURN VALUE**

20203           The *inet\_ntop()* function shall return a pointer to the buffer containing the text string if the  
20204           conversion succeeds, and NULL otherwise, and set *errno* to indicate the error.

20205           The *inet\_pton()* function shall return 1 if the conversion succeeds, with the address pointed to by  
20206           *dst* in network byte order. It shall return 0 if the input is not a valid IPv4 dotted-decimal string or  
20207           a valid IPv6 address string, or -1 with *errno* set to [EAFNOSUPPORT] if the *af* argument is  
20208           unknown.

20209 **ERRORS**

20210           The *inet\_ntop()* and *inet\_pton()* functions shall fail if:

20211           [EAFNOSUPPORT]

20212                           The *af* argument is invalid.

20213           [ENOSPC]        The size of the *inet\_ntop()* result buffer is inadequate.

20214 **EXAMPLES**

20215           None.

20216 **APPLICATION USAGE**

20217           None.

20218 **RATIONALE**

20219           None.

20220 **FUTURE DIRECTIONS**

20221           None.

20222 **SEE ALSO**

20223           The Base Definitions volume of IEEE Std. 1003.1-200x, <**arpa/inet.h**>

20224 **CHANGE HISTORY**

20225           First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

20226           Marked as part of the IPv6 option.

20227           The **restrict** keyword is added to the *inet\_ntop()* and *inet\_pton()* prototypes for alignment with  
20228           the ISO/IEC 9899:1999 standard.

20229 **NAME**20230 `initstate`, `random`, `setstate`, `srandom` — pseudorandom number functions20231 **SYNOPSIS**

```
20232 xSI #include <stdlib.h>
20233 char *initstate(unsigned seed, char *state, size_t size);
20234 long random(void);
20235 char *setstate(const char *state);
20236 void srandom(unsigned seed);
20237
```

20238 **DESCRIPTION**

20239 The `random()` function uses a non-linear additive feedback random-number generator  
 20240 employing a default state array size of 31 **long** integers to return successive pseudo-random  
 20241 numbers in the range from 0 to  $2^{31}-1$ . The period of this random-number generator is  
 20242 approximately  $16 \times (2^{31}-1)$ . The size of the state array determines the period of the random-  
 20243 number generator. Increasing the state array size increases the period.

20244 With 256 bytes of state information, the period of the random-number generator is greater than  
 20245  $2^{69}$ .

20246 Like `rand()`, `random()` produces by default a sequence of numbers that can be duplicated by  
 20247 calling `srandom()` with 1 as the seed.

20248 The `srandom()` function initializes the current state array using the value of `seed`.

20249 The `initstate()` and `setstate()` functions handle restarting and changing random-number  
 20250 generators. The `initstate()` function allows a state array, pointed to by the `state` argument, to be  
 20251 initialized for future use. The `size` argument, which specifies the size in bytes of the state array, is  
 20252 used by `initstate()` to decide what type of random-number generator to use; the larger the state  
 20253 array, the more random the numbers. Values for the amount of state information are 8, 32, 64,  
 20254 128, and 256 bytes. Other values greater than 8 bytes are rounded down to the nearest one of  
 20255 these values. If `initstate()` is called with  $8 \leq \text{size} < 32$ , then `random()` uses a simple linear  
 20256 congruential random number generator. The `seed` argument specifies a starting point for the  
 20257 random-number sequence and provides for restarting at the same point. The `initstate()` function  
 20258 shall return a pointer to the previous state information array.

20259 If `initstate()` has not been called, then `random()` behaves as though `initstate()` had been called  
 20260 with `seed=1` and `size=128`.

20261 Once a state has been initialized, `setstate()` allows switching between state arrays. The array  
 20262 defined by the `state` argument is used for further random-number generation until `initstate()` is  
 20263 called or `setstate()` is called again. The `setstate()` function shall return a pointer to the previous  
 20264 state array.

20265 **RETURN VALUE**

20266 If `initstate()` is called with `size` less than 8, it shall return NULL.

20267 The `random()` function shall return the generated pseudo-random number.

20268 The `srandom()` function shall return no value.

20269 Upon successful completion, `initstate()` and `setstate()` shall return a pointer to the previous state  
 20270 array; otherwise, a null pointer shall be returned.

20271 **ERRORS**

20272 No errors are defined.

20273 **EXAMPLES**

20274 None.

20275 **APPLICATION USAGE**

20276 After initialization, a state array can be restarted at a different point in one of two ways:

- 20277 1. The *initstate()* function can be used, with the desired seed, state array, and size of the  
20278 array.
- 20279 2. The *setstate()* function, with the desired state, can be used, followed by *srandom()* with the  
20280 desired seed. The advantage of using both of these functions is that the size of the state  
20281 array does not have to be saved once it is initialized.

20282 Although some implementations of *random()* have written messages to standard error, such  
20283 implementations do not conform to this volume of IEEE Std. 1003.1-200x.

20284 Issue 5 restores the historical behavior of this function.

20285 Threaded applications should use *rand\_r()*, *erand48()*, *nrand48()*, or *jrand48()* instead of  
20286 *random()* when an independent random number sequence in multiple threads is required.20287 **RATIONALE**

20288 None.

20289 **FUTURE DIRECTIONS**

20290 None.

20291 **SEE ALSO**20292 *drand48()*, *rand()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>`20293 **CHANGE HISTORY**

20294 First released in Issue 4, Version 2.

20295 **Issue 5**

20296 Moved from X/OPEN UNIX extension to BASE.

20297 In the DESCRIPTION, the phrase “values smaller than 8” is replaced with “values greater than  
20298 or equal to 8, or less than 32”, “*size*<8” is replaced with “ $8 \leq \textit{size} < 32$ ”, and a new first paragraph  
20299 is added to the RETURN VALUE section. A note is added to the APPLICATION USAGE  
20300 indicating that these changes restore the historical behavior of the function.

20301 **Issue 6**

20302 In the DESCRIPTION, duplicate text “For values greater than or equal to 8 . . .” is removed.

20303 **NAME**

20304           insque, remque — insert or remove an element in a queue

20305 **SYNOPSIS**

```
20306 xsi #include <search.h>
20307 void insque(void *element, void *pred);
20308 void remque(void *element);
20309
```

20310 **DESCRIPTION**

20311       The *insque()* and *remque()* functions manipulate queues built from doubly-linked lists. The  
 20312       queue can be either circular or linear. An application using *insque()* or *remque()* shall ensure it  
 20313       defines a structure in which the first two members of the structure are pointers to the same type  
 20314       of structure, and any further members are application-specific. The first member of the structure  
 20315       is a forward pointer to the next entry in the queue. The second member is a backward pointer to  
 20316       the previous entry in the queue. If the queue is linear, the queue is terminated with null  
 20317       pointers. The names of the structure and of the pointer members are not subject to any special  
 20318       restriction.

20319       The *insque()* function inserts the element pointed to by *element* into a queue immediately after  
 20320       the element pointed to by *pred*.

20321       The *remque()* function removes the element pointed to by *element* from a queue.

20322       If the queue is to be used as a linear list, invoking *insque(&element, NULL)*, where *element* is the  
 20323       initial element of the queue, shall initialize the forward and backward pointers of *element* to null  
 20324       pointers.

20325       If the queue is to be used as a circular list, the application shall ensure it initializes the forward  
 20326       pointer and the backward pointer of the initial element of the queue to the element's own  
 20327       address.

20328 **RETURN VALUE**20329           The *insque()* and *remque()* functions do not return a value.20330 **ERRORS**

20331           No errors are defined.

20332 **EXAMPLES**20333           **Creating a Linear Linked List**

20334           The following example creates a linear linked list.

```
20335 #include <search.h>
20336 ...
20337 struct myque element1;
20338 struct myque element2;
20339 char *data1 = "DATA1";
20340 char *data2 = "DATA2";
20341 ...
20342 element1.data = data1;
20343 element2.data = data2;
20344 insque (&element1, NULL);
20345 insque (&element2, &element1);
```

**20346        Creating a Circular Linked List**

20347        The following example creates a circular linked list.

```
20348 #include <search.h>
20349 ...
20350 struct myque element1;
20351 struct myque element2;

20352 char *data1 = "DATA1";
20353 char *data2 = "DATA2";
20354 ...
20355 element1.data = data1;
20356 element2.data = data2;

20357 element1.fwd = &element1;
20358 element1.bck = &element1;

20359 insque (&element2, &element1);
```

**20360        Removing an Element**

20361        The following example removes the element pointed to by *element1*.

```
20362 #include <search.h>
20363 ...
20364 struct myque element1;
20365 ...
20366 remque (&element1);
```

**20367        APPLICATION USAGE**

20368        The historical implementations of these functions described the arguments as being of type **struct qelem\*** rather than as being of type **void\*** as defined here. In those implementations, **struct qelem** was commonly defined in **<search.h>** as:

```
20371 struct qelem {
20372 struct qelem *q_forw;
20373 struct qelem *q_back;
20374 };
```

20375        Applications using these functions, however, were never able to use this structure directly since it provided no room for the actual data contained in the elements. Most applications defined structures that contained the two pointers as the initial elements and also provided space for, or pointers to, the object's data. Applications that used these functions to update more than one type of table also had the problem of specifying two or more different structures with the same name, if they literally used **struct qelem** as specified.

20381        As described here, the implementations were actually expecting a structure type where the first two members were forward and backward pointers to structures. With C compilers that didn't provide function prototypes, applications used structures as specified in the DESCRIPTION above and the compiler did what the application expected.

20385        If this method had been carried forward with an ISO C standard compiler and the historical function prototype, most applications would have to be modified to cast pointers to the structures actually used to be pointers to **struct qelem** to avoid compilation warnings. By specifying **void\*** as the argument type, applications do not need to change (unless they specifically referenced **struct qelem** and depended on it being defined in **<search.h>**).

20390 **RATIONALE**

20391           None.

20392 **FUTURE DIRECTIONS**

20393           None.

20394 **SEE ALSO**

20395           The Base Definitions volume of IEEE Std. 1003.1-200x, <search.h>

20396 **CHANGE HISTORY**

20397           First released in Issue 4, Version 2.

20398 **Issue 5**

20399           Moved from X/OPEN UNIX extension to BASE.

20400 **Issue 6**

20401           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

## 20402 NAME

20403        ioctl — control a STREAMS device (**STREAMS**)

## 20404 SYNOPSIS

20405 XSR        #include &lt;stropts.h&gt;

20406        int ioctl(int *fildev*, int *request*, ... /\* *arg* \*/);

20407

## 20408 DESCRIPTION

20409        The *ioctl()* function performs a variety of control functions on STREAMS devices. For non-  
 20410 STREAMS devices, the functions performed by this call are unspecified. The *request* argument  
 20411 and an optional third argument (with varying type) are passed to and interpreted by the  
 20412 appropriate part of the STREAM associated with *fildev*.

20413        The *fildev* argument is an open file descriptor that refers to a device.

20414        The *request* argument selects the control function to be performed and shall depend on the  
 20415 STREAMS device being addressed.

20416        The *arg* argument represents additional information that is needed by this specific STREAMS  
 20417 device to perform the requested function. The type of *arg* depends upon the particular control  
 20418 request, but it is either an integer or a pointer to a device-specific data structure.

20419        The *ioctl()* commands applicable to STREAMS, their arguments, and error conditions that apply  
 20420 to each individual command are described below.

20421        The following *ioctl()* commands, with error values indicated, are applicable to all STREAMS  
 20422 files:

20423        I\_PUSH        Pushes the module whose name is pointed to by *arg* onto the top of the  
 20424 current STREAM, just below the STREAM head. It then calls the *open()*  
 20425 function of the newly-pushed module.

20426        The *ioctl()* function with the I\_PUSH command shall fail if:

20427        [EINVAL]        Invalid module name. |

20428        [ENXIO]        Open function of new module failed. |

20429        [ENXIO]        Hangup received on *fildev*. |

20430        I\_POP        Removes the module just below the STREAM head of the STREAM pointed to  
 20431 by *fildev*. The *arg* argument should be 0 in an I\_POP request.

20432        The *ioctl()* function with the I\_POP command shall fail if:

20433        [EINVAL]        No module present in the STREAM. |

20434        [ENXIO]        Hangup received on *fildev*. |

20435        I\_LOOK        Retrieves the name of the module just below the STREAM head of the  
 20436 STREAM pointed to by *fildev*, and places it in a character string pointed to by  
 20437 *arg*. The buffer pointed to by *arg* should be at least FMNAMESZ+1 bytes long,  
 20438 where FMNAMESZ is defined in <stropts.h>.

20439        The *ioctl()* function with the I\_LOOK command shall fail if:

20440        [EINVAL]        No module present in the STREAM. |

20441        I\_FLUSH        This request flushes read and/or write queues, depending on the value of *arg*.  
 20442 Valid *arg* values are:

|       |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------|----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20443 | FLUSHR                                                                           | Flush all read queues.                                                                                                                                                                                                                                                                                                                                                                                                                |
| 20444 | FLUSHW                                                                           | Flush all write queues.                                                                                                                                                                                                                                                                                                                                                                                                               |
| 20445 | FLUSHRW                                                                          | Flush all read and all write queues.                                                                                                                                                                                                                                                                                                                                                                                                  |
| 20446 | The <i>ioctl()</i> function with the <code>L_FLUSH</code> command shall fail if: |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20447 | [EINVAL]                                                                         | Invalid <i>arg</i> value.                                                                                                                                                                                                                                                                                                                                                                                                             |
| 20448 | [EAGAIN] or [ENOSR]                                                              | Unable to allocate buffers for flush message.                                                                                                                                                                                                                                                                                                                                                                                         |
| 20449 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20450 | [ENXIO]                                                                          | Hangup received on <i>fildev</i> .                                                                                                                                                                                                                                                                                                                                                                                                    |
| 20451 | I_FLUSHBAND                                                                      | Flushes a particular band of messages. The <i>arg</i> argument points to a <b>bandinfo</b> structure. The <i>bi_flag</i> member may be one of FLUSHR, FLUSHW, or FLUSHRW as described above. The <i>bi_pri</i> member determines the priority band to be flushed.                                                                                                                                                                     |
| 20452 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20453 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20454 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20455 | I_SETSIG                                                                         | Requests that the STREAMS implementation send the SIGPOLL signal to the calling process when a particular event has occurred on the STREAM associated with <i>fildev</i> . I_SETSIG supports an asynchronous processing capability in STREAMS. The value of <i>arg</i> is a bitmask that specifies the events for which the process should be signaled. It is the bitwise-inclusive OR of any combination of the following constants: |
| 20456 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20457 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20458 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20459 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20460 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20461 | S_RDNORM                                                                         | A normal (priority band set to 0) message has arrived at the head of a STREAM head read queue. A signal shall be generated even if the message is of zero length.                                                                                                                                                                                                                                                                     |
| 20462 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20463 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20464 | S_RDBAND                                                                         | A message with a non-zero priority band has arrived at the head of a STREAM head read queue. A signal shall be generated even if the message is of zero length.                                                                                                                                                                                                                                                                       |
| 20465 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20466 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20467 | S_INPUT                                                                          | A message, other than a high-priority message, has arrived at the head of a STREAM head read queue. A signal shall be generated even if the message is of zero length.                                                                                                                                                                                                                                                                |
| 20468 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20469 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20470 | S_HIPRI                                                                          | A high-priority message is present on a STREAM head read queue. A signal shall be generated even if the message is of zero length.                                                                                                                                                                                                                                                                                                    |
| 20471 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20472 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20473 | S_OUTPUT                                                                         | The write queue for normal data (priority band 0) just below the STREAM head is no longer full. This notifies the process that there is room on the queue for sending (or writing) normal data downstream.                                                                                                                                                                                                                            |
| 20474 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20475 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20476 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20477 | S_WRNORM                                                                         | Same as S_OUTPUT.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 20478 | S_WRBAND                                                                         | The write queue for a non-zero priority band just below the STREAM head is no longer full. This notifies the process that there is room on the queue for sending (or writing) priority data downstream.                                                                                                                                                                                                                               |
| 20479 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20480 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20481 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20482 | S_MSG                                                                            | A STREAMS signal message that contains the SIGPOLL signal has reached the front of the STREAM head read queue.                                                                                                                                                                                                                                                                                                                        |
| 20483 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20484 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20485 | S_ERROR                                                                          | Notification of an error condition has reached the STREAM head.                                                                                                                                                                                                                                                                                                                                                                       |
| 20486 |                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                       |

|       |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------|----------|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20487 |          | S_HANGUP  | Notification of a hangup has reached the STREAM head.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 20488 |          | S_BANDURG | When used in conjunction with S_RDBAND, SIGURG is generated instead of SIGPOLL when a priority message reaches the front of the STREAM head read queue.                                                                                                                                                                                                                                                                                                                  |
| 20489 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20490 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20491 |          |           | If <i>arg</i> is 0, the calling process shall be unregistered and shall not receive further SIGPOLL signals for the stream associated with <i>fildev</i> .                                                                                                                                                                                                                                                                                                               |
| 20492 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20493 |          |           | Processes that wish to receive SIGPOLL signals shall ensure that they explicitly register to receive them using I_SETSIG. If several processes register to receive this signal for the same event on the same STREAM, each process shall be signaled when the event occurs.                                                                                                                                                                                              |
| 20494 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20495 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20496 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20497 |          |           | The <i>ioctl()</i> function with the I_SETSIG command shall fail if:                                                                                                                                                                                                                                                                                                                                                                                                     |
| 20498 |          | [EINVAL]  | The value of <i>arg</i> is invalid.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 20499 |          | [EINVAL]  | The value of <i>arg</i> is 0 and the calling process is not registered to receive the SIGPOLL signal.                                                                                                                                                                                                                                                                                                                                                                    |
| 20500 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20501 |          | [EAGAIN]  | There were insufficient resources to store the signal request.                                                                                                                                                                                                                                                                                                                                                                                                           |
| 20502 | I_GETSIG |           | Returns the events for which the calling process is currently registered to be sent a SIGPOLL signal. The events are returned as a bitmask in an <b>int</b> pointed to by <i>arg</i> , where the events are those specified in the description of I_SETSIG above.                                                                                                                                                                                                        |
| 20503 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20504 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20505 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20506 |          |           | The <i>ioctl()</i> function with the I_GETSIG command shall fail if:                                                                                                                                                                                                                                                                                                                                                                                                     |
| 20507 |          | [EINVAL]  | Process is not registered to receive the SIGPOLL signal.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20508 | I_FIND   |           | This request compares the names of all modules currently present in the STREAM to the name pointed to by <i>arg</i> , and returns 1 if the named module is present in the STREAM, or returns 0 if the named module is not present.                                                                                                                                                                                                                                       |
| 20509 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20510 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20511 |          |           | The <i>ioctl()</i> function with the I_FIND command shall fail if:                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20512 |          | [EINVAL]  | <i>arg</i> does not contain a valid module name.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 20513 | I_PEEK   |           | This request allows a process to retrieve the information in the first message on the STREAM head read queue without taking the message off the queue. It is analogous to <i>getmsg()</i> except that this command does not remove the message from the queue. The <i>arg</i> argument points to a <b>strpeek</b> structure.                                                                                                                                             |
| 20514 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20515 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20516 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20517 |          |           | The application shall ensure that the <i>maxlen</i> member in the <b>ctlbuf</b> and <b>databuf</b> structures is set to the number of bytes of control information and/or data information, respectively, to retrieve. The <i>flags</i> member may be marked RS_HIPRI or 0, as described by <i>getmsg()</i> . If the process sets <i>flags</i> to RS_HIPRI, for example, I_PEEK shall only look for a high-priority message on the STREAM head read queue.               |
| 20518 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20519 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20520 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20521 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20522 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20523 |          |           | I_PEEK returns 1 if a message was retrieved, and returns 0 if no message was found on the STREAM head read queue, or if the RS_HIPRI flag was set in <i>flags</i> and a high-priority message was not present on the STREAM head read queue. It does not wait for a message to arrive. On return, <b>ctlbuf</b> specifies information in the control buffer, <b>databuf</b> specifies information in the data buffer, and <i>flags</i> contains the value RS_HIPRI or 0. |
| 20524 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20525 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20526 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20527 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20528 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20529 | I_SRDOPT |           | Sets the read mode using the value of the argument <i>arg</i> . Read modes are described in <i>read()</i> . Valid <i>arg</i> flags are:                                                                                                                                                                                                                                                                                                                                  |
| 20530 |          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

|       |            |           |                                                                                               |
|-------|------------|-----------|-----------------------------------------------------------------------------------------------|
| 20531 |            | RNORM     | Byte-stream mode, the default.                                                                |
| 20532 |            | RMSGD     | Message-discard mode.                                                                         |
| 20533 |            | RMSGN     | Message-nondiscard mode.                                                                      |
| 20534 |            |           | The bitwise-inclusive OR of RMSGD and RMSGN shall return [EINVAL]. The                        |
| 20535 |            |           | bitwise-inclusive OR of RNORM and either RMSGD or RMSGN shall result in                       |
| 20536 |            |           | the other flag overriding RNORM which is the default.                                         |
| 20537 |            |           | In addition, treatment of control messages by the STREAM head may be                          |
| 20538 |            |           | changed by setting any of the following flags in <i>arg</i> :                                 |
| 20539 |            | RPROTNORM | Fail <i>read()</i> with [EBADMSG] if a message containing a                                   |
| 20540 |            |           | control part is at the front of the STREAM head read queue.                                   |
| 20541 |            | RPROTDAT  | Deliver the control part of a message as data when a                                          |
| 20542 |            |           | process issues a <i>read()</i> .                                                              |
| 20543 |            | RPROTDIS  | Discard the control part of a message, delivering any data                                    |
| 20544 |            |           | portion, when a process issues a <i>read()</i> .                                              |
| 20545 |            |           | The <i>ioctl()</i> function with the I_SRDOPT command shall fail if:                          |
| 20546 |            | [EINVAL]  | The <i>arg</i> argument is not valid.                                                         |
| 20547 | I_GRDOPT   |           | Returns the current read mode setting as, described above, in an <b>int</b> pointed to        |
| 20548 |            |           | by the argument <i>arg</i> . Read modes are described in <i>read()</i> .                      |
| 20549 | I_NREAD    |           | Counts the number of data bytes in the data part of the first message on the                  |
| 20550 |            |           | STREAM head read queue and places this value in the <b>int</b> pointed to by <i>arg</i> .     |
| 20551 |            |           | The return value for the command is the number of messages on the STREAM                      |
| 20552 |            |           | head read queue. For example, if 0 is returned in <i>arg</i> , but the <i>ioctl()</i> return  |
| 20553 |            |           | value is greater than 0, this indicates that a zero-length message is next on the             |
| 20554 |            |           | queue.                                                                                        |
| 20555 | I_FDINSERT |           | Creates a message from specified buffer(s), adds information about another                    |
| 20556 |            |           | STREAM, and sends the message downstream. The message contains a                              |
| 20557 |            |           | control part and an optional data part. The data and control parts to be sent                 |
| 20558 |            |           | are distinguished by placement in separate buffers, as described below. The                   |
| 20559 |            |           | <i>arg</i> argument points to a <b>strfdinsert</b> structure.                                 |
| 20560 |            |           | The application shall ensure that the <i>len</i> member in the <b>ctlbuf strbuf</b> structure |
| 20561 |            |           | is set to the size of a <b>t_uscalar_t</b> plus the number of bytes of control                |
| 20562 |            |           | information to be sent with the message. The <i>fildev</i> member specifies the file          |
| 20563 |            |           | descriptor of the other STREAM, and the <i>offset</i> member, which must be                   |
| 20564 |            |           | suitably aligned for use as a <b>t_uscalar_t</b> , specifies the offset from the start of     |
| 20565 |            |           | the control buffer where I_FDINSERT shall store a <b>t_uscalar_t</b> whose                    |
| 20566 |            |           | interpretation is specific to the STREAM end. The application shall ensure that               |
| 20567 |            |           | the <i>len</i> member in the <b>datbuf strbuf</b> structure is set to the number of bytes of  |
| 20568 |            |           | data information to be sent with the message, or to 0 if no data part is to be                |
| 20569 |            |           | sent.                                                                                         |
| 20570 |            |           | The <i>flags</i> member specifies the type of message to be created. A normal                 |
| 20571 |            |           | message is created if <i>flags</i> is set to 0, and a high-priority message is created if     |
| 20572 |            |           | <i>flags</i> is set to RS_HIPRI. For non-priority messages, I_FDINSERT shall block if         |
| 20573 |            |           | the STREAM write queue is full due to internal flow control conditions. For                   |
| 20574 |            |           | priority messages, I_FDINSERT does not block on this condition. For non-                      |
| 20575 |            |           | priority messages, I_FDINSERT does not block when the write queue is full                     |

20576 and O\_NONBLOCK is set. Instead, it fails and sets *errno* to [EAGAIN].

20577 I\_FDINSERT also blocks, unless prevented by lack of internal resources,  
 20578 waiting for the availability of message blocks in the STREAM, regardless of  
 20579 priority or whether O\_NONBLOCK has been specified. No partial message is  
 20580 sent.

20581 The *ioctl()* function with the I\_FDINSERT command shall fail if:

20582 [EAGAIN] A non-priority message is specified, the O\_NONBLOCK  
 20583 flag is set, and the STREAM write queue is full due to  
 20584 internal flow control conditions.

20585 [EAGAIN] or [ENOSR]  
 20586 Buffers cannot be allocated for the message that is to be  
 20587 created.

20588 [EINVAL] One of the following:

20589 — The *fildev* member of the **strfdinsert** structure is not a  
 20590 valid, open STREAM file descriptor.

20591 — The size of a **t\_uscalar\_t** plus *offset* is greater than the *len*  
 20592 member for the buffer specified through **ctlbuf**.

20593 — The *offset* member does not specify a properly-aligned  
 20594 location in the data buffer.

20595 — An undefined value is stored in *flags*.

20596 [ENXIO] Hangupt received on the STREAM identified by either the  
 20597 *fildev* argument or the *fildev* member of the **strfdinsert**  
 20598 structure.

20599 [ERANGE] The *len* member for the buffer specified through **databuf**  
 20600 does not fall within the range specified by the maximum  
 20601 and minimum packet sizes of the topmost STREAM module  
 20602 or the *len* member for the buffer specified through **databuf**  
 20603 is larger than the maximum configured size of the data part  
 20604 of a message; or the *len* member for the buffer specified  
 20605 through **ctlbuf** is larger than the maximum configured size  
 20606 of the control part of a message.

20607 I\_STR Constructs an internal STREAMS *ioctl()* message from the data pointed to by  
 20608 *arg*, and sends that message downstream.

20609 This mechanism is provided to send *ioctl()* requests to downstream modules  
 20610 and drivers. It allows information to be sent with *ioctl()*, and returns to the  
 20611 process any information sent upstream by the downstream recipient. I\_STR  
 20612 blocks until the system responds with either a positive or negative  
 20613 acknowledgement message, or until the request times out after some period of  
 20614 time. If the request times out, it fails with *errno* set to [ETIME].

20615 At most, one I\_STR can be active on a STREAM. Further I\_STR calls shall  
 20616 block until the active I\_STR completes at the STREAM head. The default  
 20617 timeout interval for these requests is 15 seconds. The O\_NONBLOCK flag has  
 20618 no effect on this call.

20619 To send requests downstream, the application shall ensure that *arg* points to a  
 20620 **strioc** structure.

|       |          |                                                                                                                                                                                                                                         |
|-------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20621 |          | The <i>ic_cmd</i> member is the internal <i>ioctl()</i> command intended for a downstream module or driver and <i>ic_timeout</i> is the number of seconds (−1=infinite, 0=use implementation-defined timeout interval, >0=as specified) |
| 20622 |          | an I_STR request shall wait for acknowledgement before timing out. <i>ic_len</i> is                                                                                                                                                     |
| 20623 |          | the number of bytes in the data argument, and <i>ic_dp</i> is a pointer to the data                                                                                                                                                     |
| 20624 |          | argument. The <i>ic_len</i> member has two uses: on input, it contains the length of                                                                                                                                                    |
| 20625 |          | the data argument passed in, and on return from the command, it contains the                                                                                                                                                            |
| 20626 |          | number of bytes being returned to the process (the buffer pointed to by <i>ic_dp</i>                                                                                                                                                    |
| 20627 |          | should be large enough to contain the maximum amount of data that any                                                                                                                                                                   |
| 20628 |          | module or the driver in the STREAM can return).                                                                                                                                                                                         |
| 20629 |          |                                                                                                                                                                                                                                         |
| 20630 |          |                                                                                                                                                                                                                                         |
| 20631 |          | The STREAM head shall convert the information pointed to by the <b>strioc</b>                                                                                                                                                           |
| 20632 |          | structure to an internal <i>ioctl()</i> command message and sends it downstream.                                                                                                                                                        |
| 20633 |          | The <i>ioctl()</i> function with the I_STR command shall fail if:                                                                                                                                                                       |
| 20634 |          | [EAGAIN] or [ENOSR]                                                                                                                                                                                                                     |
| 20635 |          | Unable to allocate buffers for the <i>ioctl()</i> message.                                                                                                                                                                              |
| 20636 |          | [EINVAL] The <i>ic_len</i> member is less than 0 or larger than the                                                                                                                                                                     |
| 20637 |          | maximum configured size of the data part of a message, or                                                                                                                                                                               |
| 20638 |          | <i>ic_timeout</i> is less than −1.                                                                                                                                                                                                      |
| 20639 |          | [ENXIO] Hangup received on <i>fil</i> des.                                                                                                                                                                                              |
| 20640 |          | [ETIME] A downstream <i>ioctl()</i> timed out before acknowledgement                                                                                                                                                                    |
| 20641 |          | was received.                                                                                                                                                                                                                           |
| 20642 |          | An I_STR can also fail while waiting for an acknowledgement if a message                                                                                                                                                                |
| 20643 |          | indicating an error or a hangup is received at the STREAM head. In addition,                                                                                                                                                            |
| 20644 |          | an error code can be returned in the positive or negative acknowledgement                                                                                                                                                               |
| 20645 |          | message, in the event the <i>ioctl()</i> command sent downstream fails. For these                                                                                                                                                       |
| 20646 |          | cases, I_STR fails with <i>errno</i> set to the value in the message.                                                                                                                                                                   |
| 20647 | I_SWROPT | Sets the write mode using the value of the argument <i>arg</i> . Valid bit settings for                                                                                                                                                 |
| 20648 |          | <i>arg</i> are:                                                                                                                                                                                                                         |
| 20649 |          | SNDZERO Send a zero-length message downstream when a <i>write()</i> of                                                                                                                                                                  |
| 20650 |          | 0 bytes occurs. To not send a zero-length message when a                                                                                                                                                                                |
| 20651 |          | <i>write()</i> of 0 bytes occurs, the application shall ensure that                                                                                                                                                                     |
| 20652 |          | this bit is not set in <i>arg</i> (for example, <i>arg</i> would be set to 0).                                                                                                                                                          |
| 20653 |          | The <i>ioctl()</i> function with the I_SWROPT command shall fail if:                                                                                                                                                                    |
| 20654 |          | [EINVAL] <i>arg</i> is not the above value.                                                                                                                                                                                             |
| 20655 | I_GWROPT | Returns the current write mode setting, as described above, in the <b>int</b> that is                                                                                                                                                   |
| 20656 |          | pointed to by the argument <i>arg</i> .                                                                                                                                                                                                 |
| 20657 | I_SENDFD | I_SENDFD creates a new reference to the open file description associated with                                                                                                                                                           |
| 20658 |          | the file descriptor <i>arg</i> , and writes a message on the STREAMS-based pipe                                                                                                                                                         |
| 20659 |          | <i>fil</i> des containing this reference, together with the user ID and group ID of the                                                                                                                                                 |
| 20660 |          | calling process.                                                                                                                                                                                                                        |
| 20661 |          | The <i>ioctl()</i> function with the I_SENDFD command shall fail if:                                                                                                                                                                    |
| 20662 |          | [EAGAIN] The sending STREAM is unable to allocate a message block                                                                                                                                                                       |
| 20663 |          | to contain the file pointer; or the read queue of the receiving                                                                                                                                                                         |
| 20664 |          | STREAM head is full and cannot accept the message sent by                                                                                                                                                                               |
| 20665 |          | I_SENDFD.                                                                                                                                                                                                                               |

|       |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------|----------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20666 |          | [EBADF]             | The <i>arg</i> argument is not a valid, open file descriptor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 20667 |          | [EINVAL]            | The <i>fildev</i> argument is not connected to a STREAM pipe.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 20668 |          | [ENXIO]             | Hangup received on <i>fildev</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 20669 | I_RECVFD |                     | Retrieves the reference to an open file description from a message written to a STREAMS-based pipe using the I_SENDFD command, and allocates a new file descriptor in the calling process that refers to this open file description. The <i>arg</i> argument is a pointer to a <b>strrecvfd</b> data structure as defined in <b>&lt;stropts.h&gt;</b> .                                                                                                                                                                                                                                                                                                         |
| 20670 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20671 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20672 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20673 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20674 |          |                     | The <i>fd</i> member is a file descriptor. The <i>uid</i> and <i>gid</i> members are the effective user ID and effective group ID, respectively, of the sending process.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 20675 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20676 |          |                     | If O_NONBLOCK is not set, I_RECVFD blocks until a message is present at the STREAM head. If O_NONBLOCK is set, I_RECVFD fails with <i>errno</i> set to [EAGAIN] if no message is present at the STREAM head.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 20677 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20678 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20679 |          |                     | If the message at the STREAM head is a message sent by an I_SENDFD, a new file descriptor is allocated for the open file descriptor referenced in the message. The new file descriptor is placed in the <i>fd</i> member of the <b>strrecvfd</b> structure pointed to by <i>arg</i> .                                                                                                                                                                                                                                                                                                                                                                           |
| 20680 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20681 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20682 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20683 |          |                     | The <i>ioctl()</i> function with the I_RECVFD command shall fail if:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 20684 |          | [EAGAIN]            | A message is not present at the STREAM head read queue and the O_NONBLOCK flag is set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 20685 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20686 |          | [EBADMSG]           | The message at the STREAM head read queue is not a message containing a passed file descriptor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20687 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20688 |          | [EMFILE]            | The process has the maximum number of file descriptors currently open that it is allowed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 20689 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20690 |          | [ENXIO]             | Hangup received on <i>fildev</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 20691 | I_LIST   |                     | This request allows the process to list all the module names on the STREAM, up to and including the topmost driver name. If <i>arg</i> is a null pointer, the return value is the number of modules, including the driver, that are on the STREAM pointed to by <i>fildev</i> . This lets the process allocate enough space for the module names. Otherwise, it should point to a <b>str_list</b> structure.                                                                                                                                                                                                                                                    |
| 20692 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20693 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20694 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20695 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20696 |          |                     | The <i>sl_nmods</i> member indicates the number of entries the process has allocated in the array. Upon return, the <i>sl_modlist</i> member of the <b>str_list</b> structure contains the list of module names, and the number of entries that have been filled into the <i>sl_modlist</i> array is found in the <i>sl_nmods</i> member (the number includes the number of modules including the driver). The return value from <i>ioctl()</i> is 0. The entries are filled in starting at the top of the STREAM and continuing downstream until either the end of the STREAM is reached, or the number of requested modules ( <i>sl_nmods</i> ) is satisfied. |
| 20697 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20698 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20699 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20700 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20701 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20702 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20703 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20704 |          |                     | The <i>ioctl()</i> function with the I_LIST command shall fail if:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 20705 |          | [EINVAL]            | The <i>sl_nmods</i> member is less than 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 20706 |          | [EAGAIN] or [ENOSR] |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 20707 |          |                     | Unable to allocate buffers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 20708 | I_ATMARK |                     | This request allows the process to see if the message at the head of the STREAM head read queue is marked by some module downstream. The <i>arg</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 20709 |          |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

|       |             |                                                                                           |
|-------|-------------|-------------------------------------------------------------------------------------------|
| 20710 |             | argument determines how the checking is done when there may be multiple                   |
| 20711 |             | marked messages on the STREAM head read queue. It may take on the                         |
| 20712 |             | following values:                                                                         |
| 20713 |             | ANYMARK      Check if the message is marked.                                              |
| 20714 |             | LASTMARK     Check if the message is the last one marked on the queue.                    |
| 20715 |             | The bitwise-inclusive OR of the flags ANYMARK and LASTMARK is                             |
| 20716 |             | permitted.                                                                                |
| 20717 |             | The return value is 1 if the mark condition is satisfied; otherwise, the value is         |
| 20718 |             | 0.                                                                                        |
| 20719 |             | The <i>ioctl()</i> function with the L_ATMARK command shall fail if:                      |
| 20720 |             | [EINVAL]      Invalid <i>arg</i> value.                                                   |
| 20721 | I_CKBAND    | Check if the message of a given priority band exists on the STREAM head                   |
| 20722 |             | read queue. This returns 1 if a message of the given priority exists, 0 if no such        |
| 20723 |             | message exists, or -1 on error. <i>arg</i> should be of type <b>int</b> .                 |
| 20724 |             | The <i>ioctl()</i> function with the L_CKBAND command shall fail if:                      |
| 20725 |             | [EINVAL]      Invalid <i>arg</i> value.                                                   |
| 20726 | I_GETBAND   | Return the priority band of the first message on the STREAM head read queue               |
| 20727 |             | in the integer referenced by <i>arg</i> .                                                 |
| 20728 |             | The <i>ioctl()</i> function with the L_GETBAND command shall fail if:                     |
| 20729 |             | [ENODATA]     No message on the STREAM head read queue.                                   |
| 20730 | I_CANPUT    | Check if a certain band is writable. <i>arg</i> is set to the priority band in question.  |
| 20731 |             | The return value is 0 if the band is flow-controlled, 1 if the band is writable, or       |
| 20732 |             | -1 on error.                                                                              |
| 20733 |             | The <i>ioctl()</i> function with the L_CANPUT command shall fail if:                      |
| 20734 |             | [EINVAL]      Invalid <i>arg</i> value.                                                   |
| 20735 | I_SETCLTIME | This request allows the process to set the time the STREAM head shall delay               |
| 20736 |             | when a STREAM is closing and there is data on the write queues. Before                    |
| 20737 |             | closing each module or driver, if there is data on its write queue, the STREAM            |
| 20738 |             | head shall delay for the specified amount of time to allow the data to drain. If,         |
| 20739 |             | after the delay, data is still present, it shall be flushed. The <i>arg</i> argument is a |
| 20740 |             | pointer to an integer specifying the number of milliseconds to delay, rounded             |
| 20741 |             | up to the nearest valid value. If I_SETCLTIME is not performed on a STREAM,               |
| 20742 |             | an implementation-defined default timeout interval is used.                               |
| 20743 |             | The <i>ioctl()</i> function with the L_SETCLTIME command shall fail if:                   |
| 20744 |             | [EINVAL]      Invalid <i>arg</i> value.                                                   |
| 20745 | I_GETCLTIME | This request returns the close time delay in the integer pointed to by <i>arg</i> .       |

20746 **Multiplexed STREAMS Configurations**

20747 The following commands are used for connecting and disconnecting multiplexed STREAMS  
20748 configurations. These commands use an implementation-defined default timeout interval.

20749 **I\_LINK** Connects two STREAMs, where *fildev* is the file descriptor of the STREAM  
20750 connected to the multiplexing driver, and *arg* is the file descriptor of the  
20751 STREAM connected to another driver. The STREAM designated by *arg* is  
20752 connected below the multiplexing driver. I\_LINK requires the multiplexing  
20753 driver to send an acknowledgement message to the STREAM head regarding  
20754 the connection. This call returns a multiplexer ID number (an identifier used  
20755 to disconnect the multiplexer; see I\_UNLINK) on success, and -1 on failure.

20756 The *ioctl()* function with the I\_LINK command shall fail if:

20757 [ENXIO] Hangup received on *fildev*.

20758 [ETIME] Timeout before acknowledgement message was received at  
20759 STREAM head.

20760 [EAGAIN] or [ENOSR]

20761 Unable to allocate STREAMS storage to perform the  
20762 I\_LINK.

20763 [EBADF] The *arg* argument is not a valid, open file descriptor.

20764 [EINVAL] The *fildev* argument does not support multiplexing; or *arg* is  
20765 not a STREAM or is already connected downstream from a  
20766 multiplexer; or the specified I\_LINK operation would  
20767 connect the STREAM head in more than one place in the  
20768 multiplexed STREAM.

20769 An I\_LINK can also fail while waiting for the multiplexing driver to  
20770 acknowledge the request, if a message indicating an error or a hangup is  
20771 received at the STREAM head of *fildev*. In addition, an error code can be  
20772 returned in the positive or negative acknowledgement message. For these  
20773 cases, I\_LINK fails with *errno* set to the value in the message.

20774 **I\_UNLINK** Disconnects the two STREAMs specified by *fildev* and *arg*. *fildev* is the file  
20775 descriptor of the STREAM connected to the multiplexing driver. The *arg*  
20776 argument is the multiplexer ID number that was returned by the I\_LINK  
20777 *ioctl()* command when a STREAM was connected downstream from the  
20778 multiplexing driver. If *arg* is MUXID\_ALL, then all STREAMs that were  
20779 connected to *fildev* are disconnected. As in I\_LINK, this command requires  
20780 acknowledgement.

20781 The *ioctl()* function with the I\_UNLINK command shall fail if:

20782 [ENXIO] Hangup received on *fildev*.

20783 [ETIME] Timeout before acknowledgement message was received at  
20784 STREAM head.

20785 [EAGAIN] or [ENOSR]

20786 Unable to allocate buffers for the acknowledgement  
20787 message.

20788 [EINVAL] Invalid multiplexer ID number.

20789 An I\_UNLINK can also fail while waiting for the multiplexing driver to  
 20790 acknowledge the request if a message indicating an error or a hangup is  
 20791 received at the STREAM head of *filde*s. In addition, an error code can be  
 20792 returned in the positive or negative acknowledgement message. For these  
 20793 cases, I\_UNLINK fails with *errno* set to the value in the message.

20794 I\_PLINK Creates a *persistent connection* between two STREAMs, where *filde*s is the file  
 20795 descriptor of the STREAM connected to the multiplexing driver, and *arg* is the  
 20796 file descriptor of the STREAM connected to another driver. This call creates a  
 20797 persistent connection which can exist even if the file descriptor *filde*s  
 20798 associated with the upper STREAM to the multiplexing driver is closed. The  
 20799 STREAM designated by *arg* gets connected via a persistent connection below  
 20800 the multiplexing driver. I\_PLINK requires the multiplexing driver to send an  
 20801 acknowledgement message to the STREAM head. This call returns a  
 20802 multiplexer ID number (an identifier that may be used to disconnect the  
 20803 multiplexer; see I\_PUNLINK) on success, and -1 on failure.

20804 The *ioctl()* function with the I\_PLINK command shall fail if:

20805 [ENXIO] Hangup received on *filde*s.

20806 [ETIME] Timeout before acknowledgement message was received at  
 20807 STREAM head.

20808 [EAGAIN] or [ENOSR] Unable to allocate STREAMS storage to perform the  
 20809 I\_PLINK.  
 20810

20811 [EBADF] The *arg* argument is not a valid, open file descriptor.

20812 [EINVAL] The *filde*s argument does not support multiplexing; or *arg* is  
 20813 not a STREAM or is already connected downstream from a  
 20814 multiplexer; or the specified I\_PLINK operation would  
 20815 connect the STREAM head in more than one place in the  
 20816 multiplexed STREAM.

20817 An I\_PLINK can also fail while waiting for the multiplexing driver to  
 20818 acknowledge the request, if a message indicating an error or a hangup is  
 20819 received at the STREAM head of *filde*s. In addition, an error code can be  
 20820 returned in the positive or negative acknowledgement message. For these  
 20821 cases, I\_PLINK fails with *errno* set to the value in the message.

20822 I\_PUNLINK Disconnects the two STREAMs specified by *filde*s and *arg* from a persistent  
 20823 connection. The *filde*s argument is the file descriptor of the STREAM  
 20824 connected to the multiplexing driver. The *arg* argument is the multiplexer ID  
 20825 number that was returned by the I\_PLINK *ioctl()* command when a STREAM  
 20826 was connected downstream from the multiplexing driver. If *arg* is  
 20827 MUXID\_ALL, then all STREAMs which are persistent connections to *filde*s  
 20828 are disconnected. As in I\_PLINK, this command requires the multiplexing driver  
 20829 to acknowledge the request.

20830 The *ioctl()* function with the I\_PUNLINK command shall fail if:

20831 [ENXIO] Hangup received on *filde*s.

20832 [ETIME] Timeout before acknowledgement message was received at  
 20833 STREAM head.

20834 [EAGAIN] or [ENOSR]  
 20835 Unable to allocate buffers for the acknowledgement  
 20836 message.

20837 [EINVAL] Invalid multiplexer ID number.

20838 An I\_PUNLINK can also fail while waiting for the multiplexing driver to  
 20839 acknowledge the request if a message indicating an error or a hangup is  
 20840 received at the STREAM head of *fildev*. In addition, an error code can be  
 20841 returned in the positive or negative acknowledgement message. For these  
 20842 cases, I\_PUNLINK fails with *errno* set to the value in the message.

#### 20843 RETURN VALUE

20844 Upon successful completion, *ioctl()* shall return a value other than  $-1$  that depends upon the  
 20845 STREAMS device control function. Otherwise, it shall return  $-1$  and set *errno* to indicate the  
 20846 error.

#### 20847 ERRORS

20848 Under the following general conditions, *ioctl()* shall fail if:

20849 [EBADF] The *fildev* argument is not a valid open file descriptor.

20850 [EINTR] A signal was caught during the *ioctl()* operation.

20851 [EINVAL] The STREAM or multiplexer referenced by *fildev* is linked (directly or  
 20852 indirectly) downstream from a multiplexer.

20853 If an underlying device driver detects an error, then *ioctl()* shall fail if:

20854 [EINVAL] The *request* or *arg* argument is not valid for this device.

20855 [EIO] Some physical I/O error has occurred.

20856 [ENOTTY] The *fildev* argument is not associated with a STREAMS device that accepts  
 20857 control functions.

20858 [ENXIO] The *request* and *arg* arguments are valid for this device driver, but the service  
 20859 requested cannot be performed on this particular sub-device.

20860 [ENODEV] The *fildev* argument refers to a valid STREAMS device, but the corresponding  
 20861 device driver does not support the *ioctl()* function.

20862 If a STREAM is connected downstream from a multiplexer, any *ioctl()* command except  
 20863 I\_UNLINK and I\_PUNLINK shall set *errno* to [EINVAL].

#### 20864 EXAMPLES

20865 None.

#### 20866 APPLICATION USAGE

20867 The implementation-defined timeout interval for STREAMS has historically been 15 seconds.

#### 20868 RATIONALE

20869 None.

#### 20870 FUTURE DIRECTIONS

20871 None.

#### 20872 SEE ALSO

20873 *close()*, *fcntl()*, *getmsg()*, *open()*, *pipe()*, *poll()*, *putmsg()*, *read()*, *sigaction()*, *write()*, the Base  
 20874 Definitions volume of IEEE Std. 1003.1-200x, <*stropts.h*>, Section 2.6 (on page 539)

20875 **CHANGE HISTORY**

20876 First released in Issue 4, Version 2.

20877 **Issue 5**

20878 Moved from X/OPEN UNIX extension to BASE.

20879 **Issue 6**

20880 The Open Group corrigenda item U028/4 has been applied, correcting text in the I\_FDINSERT, [EINVAL] case to refer to *ctlbuf*.

20882 This function is marked as part of the XSI STREAMS Option Group.

20883 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

20884 **NAME**

20885           isalnum — test for an alphanumeric character

20886 **SYNOPSIS**

20887           #include <ctype.h>

20888           int isalnum(int c);

20889 **DESCRIPTION**

20890 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
20891           conflict between the requirements described here and the ISO C standard is unintentional. This  
20892           volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

20893           The *isalnum()* function tests whether *c* is a character of class **alpha** or **digit** in the program's  
20894           current locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

20895           In all cases *c* is an **int**, the value of which the application shall ensure is representable as an  
20896           **unsigned char** or equal to the value of the macro EOF. If the argument has any other value, the  
20897           behavior is undefined.

20898 **RETURN VALUE**

20899           The *isalnum()* function shall return non-zero if *c* is an alphanumeric character; otherwise, it shall  
20900           return 0.

20901 **ERRORS**

20902           No errors are defined.

20903 **EXAMPLES**

20904           None.

20905 **APPLICATION USAGE**

20906           To ensure applications portability, especially across natural languages, only this function and  
20907           those listed in the SEE ALSO section should be used for character classification.

20908 **RATIONALE**

20909           None.

20910 **FUTURE DIRECTIONS**

20911           None.

20912 **SEE ALSO**

20913           *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*,  
20914           *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, <stdio.h>, the Base  
20915           Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

20916 **CHANGE HISTORY**

20917           First released in Issue 1. Derived from Issue 1 of the SVID.

20918 **Issue 4**

20919           The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
20920           functional differences between this issue and Issue 3. Operation in the C locale is no longer  
20921           described explicitly on this reference page.

20922 **Issue 6**

20923           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

20924 **NAME**

20925 isalpha — test for an alphabetic character

20926 **SYNOPSIS**

20927 #include <ctype.h>

20928 int isalpha(int c);

20929 **DESCRIPTION**

20930 **cx** The functionality described on this reference page is aligned with the ISO C standard. Any  
20931 conflict between the requirements described here and the ISO C standard is unintentional. This  
20932 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

20933 The *isalpha()* function shall test whether *c* is a character of class **alpha** in the program's current  
20934 locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

20935 In all cases *c* is an **int**, the value of which the application shall ensure is representable as an  
20936 **unsigned char** or equal to the value of the macro EOF. If the argument has any other value, the  
20937 behavior is undefined.

20938 **RETURN VALUE**

20939 The *isalpha()* function shall return non-zero if *c* is an alphabetic character; otherwise, it shall  
20940 return 0.

20941 **ERRORS**

20942 No errors are defined.

20943 **EXAMPLES**

20944 None.

20945 **APPLICATION USAGE**

20946 To ensure applications portability, especially across natural languages, only this function and  
20947 those listed in the SEE ALSO section should be used for character classification.

20948 **RATIONALE**

20949 None.

20950 **FUTURE DIRECTIONS**

20951 None.

20952 **SEE ALSO**

20953 *isalnum()*, *isctrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
20954 *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, <stdio.h>,  
20955 the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

20956 **CHANGE HISTORY**

20957 First released in Issue 1. Derived from Issue 1 of the SVID.

20958 **Issue 4**

20959 The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
20960 functional differences between this issue and Issue 3. Operation in the C locale is no longer  
20961 described explicitly on this reference page.

20962 **Issue 6**

20963 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

20964 **NAME**

20965           isascii — test for a 7-bit US-ASCII character

20966 **SYNOPSIS**

20967 xSI       #include &lt;ctype.h&gt;

20968           int isascii(int c);

20969

20970 **DESCRIPTION**20971           The *isascii()* function tests whether *c* is a 7-bit US-ASCII character code.20972           The *isascii()* function is defined on all integer values.20973 **RETURN VALUE**20974           The *isascii()* function shall return non-zero if *c* is a 7-bit US-ASCII character code between 0 and

20975           octal 0177 inclusive; otherwise, it shall return 0.

20976 **ERRORS**

20977           No errors are defined.

20978 **EXAMPLES**

20979           None.

20980 **APPLICATION USAGE**

20981           None.

20982 **RATIONALE**

20983           None.

20984 **FUTURE DIRECTIONS**

20985           None.

20986 **SEE ALSO**

20987           The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;ctype.h&gt;

20988 **CHANGE HISTORY**

20989           First released in Issue 1. Derived from Issue 1 of the SVID.

20990 **NAME**

20991 isastream — test a file descriptor (**STREAMS**)

20992 **SYNOPSIS**

```
20993 xSR #include <stropts.h>
```

```
20994 int isastream(int fildev);
```

20995

20996 **DESCRIPTION**

20997 The *isastream()* function shall test whether *fildev*, an open file descriptor, is associated with a  
20998 STREAMS-based file.

20999 **RETURN VALUE**

21000 Upon successful completion, *isastream()* shall return 1 if *fildev* refers to a STREAMS-based file  
21001 and 0 if not. Otherwise, *isastream()* shall return -1 and set *errno* to indicate the error.

21002 **ERRORS**

21003 The *isastream()* function shall fail if:

21004 [EBADF] The *fildev* argument is not a valid open file descriptor.

21005 **EXAMPLES**

21006 None.

21007 **APPLICATION USAGE**

21008 None.

21009 **RATIONALE**

21010 None.

21011 **FUTURE DIRECTIONS**

21012 None.

21013 **SEE ALSO**

21014 The Base Definitions volume of IEEE Std. 1003.1-200x, <stropts.h>

21015 **CHANGE HISTORY**

21016 First released in Issue 4, Version 2.

21017 **Issue 5**

21018 Moved from X/OPEN UNIX extension to BASE.

21019 **NAME**

21020           isatty — test for a terminal device

21021 **SYNOPSIS**

21022           #include <unistd.h>

21023           int isatty(int *fildev*);

21024 **DESCRIPTION**

21025           The *isatty()* function shall test whether *fildev*, an open file descriptor, is associated with a  
21026           terminal device.

21027 **RETURN VALUE**

21028           The *isatty()* function shall return 1 if *fildev* is associated with a terminal; otherwise, it shall return  
21029           0 and may set *errno* to indicate the error.

21030 **ERRORS**

21031           The *isatty()* function may fail if:

21032           [EBADF]           The *fildev* argument is not a valid open file descriptor.

21033           [ENOTTY]         The *fildev* argument is not associated with a terminal.

21034 **EXAMPLES**

21035           None.

21036 **APPLICATION USAGE**

21037           The *isatty()* function does not necessarily indicate that a human being is available for interaction  
21038           via *fildev*. It is quite possible that non-terminal devices are connected to the communications  
21039           line.

21040 **RATIONALE**

21041           None.

21042 **FUTURE DIRECTIONS**

21043           None.

21044 **SEE ALSO**

21045           The Base Definitions volume of IEEE Std. 1003.1-200x, <unistd.h>

21046 **CHANGE HISTORY**

21047           First released in Issue 1. Derived from Issue 1 of the SVID.

21048 **Issue 4**

21049           The <unistd.h> header is added to the SYNOPSIS section.

21050           In the RETURN VALUE section, the sentence indicating that this function may set *errno* is  
21051           marked as an extension.

21052           The errors [EBADF] and [ENOTTY] are marked as extensions.

21053 **Issue 6**

21054           The following new requirements on POSIX implementations derive from alignment with the  
21055           Single UNIX Specification:

- 21056           • The optional setting of *errno* to indicate an error is added.
- 21057           • The [EBADF] and [ENOTTY] optional error conditions are added.

21058 **NAME**

21059           isblank — test for a blank character

21060 **SYNOPSIS**

21061           #include <ctype.h>

21062           int isblank(int c);

21063 **DESCRIPTION**

21064 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
21065       conflict between the requirements described here and the ISO C standard is unintentional. This  
21066       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21067       The *isblank()* function shall test whether *c* is a character of class **blank** in the program's current  
21068       locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale. In all cases *c*  
21069       is a type **int**, the value of which the application shall ensure is a character representable as an  
21070       **unsigned char** or equal to the value of the macro EOF. If the argument has any other value, the  
21071       behavior is undefined.

21072 **RETURN VALUE**

21073       The *isblank()* function shall return non-zero if *c* is a <blank> character; otherwise, it shall return  
21074       0.

21075 **ERRORS**

21076       No errors are defined.

21077 **EXAMPLES**

21078       None.

21079 **APPLICATION USAGE**

21080       To ensure applications portability, especially across natural languages, only this function and  
21081       those listed in the SEE ALSO section should be used for character classification.

21082 **RATIONALE**

21083       None.

21084 **FUTURE DIRECTIONS**

21085       None.

21086 **SEE ALSO**

21087       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
21088       *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>

21089 **CHANGE HISTORY**

21090       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21091 **NAME**

21092 isctrl — test for a control character

21093 **SYNOPSIS**

21094 #include &lt;ctype.h&gt;

21095 int isctrl(int c);

21096 **DESCRIPTION**

21097 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
21098 conflict between the requirements described here and the ISO C standard is unintentional. This  
21099 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21100 The *isctrl()* function shall test whether *c* is a character of class **cntrl** in the program's current  
21101 locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

21102 In all cases *c* is a type **int**, the value of which the application shall ensure is a character  
21103 representable as an **unsigned char** or equal to the value of the macro EOF. If the argument has  
21104 any other value, the behavior is undefined.

21105 **RETURN VALUE**21106 The *isctrl()* function shall return non-zero if *c* is a control character; otherwise, it shall return 0.21107 **ERRORS**

21108 No errors are defined.

21109 **EXAMPLES**

21110 None.

21111 **APPLICATION USAGE**

21112 To ensure applications portability, especially across natural languages, only this function and  
21113 those listed in the SEE ALSO section should be used for character classification.

21114 **RATIONALE**

21115 None.

21116 **FUTURE DIRECTIONS**

21117 None.

21118 **SEE ALSO**

21119 *isalnum()*, *isalpha()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
21120 *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base  
21121 Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

21122 **CHANGE HISTORY**

21123 First released in Issue 1. Derived from Issue 1 of the SVID.

21124 **Issue 4**

21125 The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
21126 functional differences between this issue and Issue 3. Operation in the C locale is no longer  
21127 described explicitly on this reference page.

21128 **Issue 6**

21129 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21130 **NAME**

21131            isdigit — test for a decimal digit

21132 **SYNOPSIS**

21133            #include <ctype.h>

21134            int isdigit(int c);

21135 **DESCRIPTION**

21136 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
21137            conflict between the requirements described here and the ISO C standard is unintentional. This  
21138            volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21139            The *isdigit()* function shall test whether *c* is a character of class **digit** in the program's current  
21140            locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

21141            In all cases *c* is an **int**, the value of which the application shall ensure is a character representable  
21142            as an **unsigned char** or equal to the value of the macro EOF. If the argument has any other value,  
21143            the behavior is undefined.

21144 **RETURN VALUE**

21145            The *isdigit()* function shall return non-zero if *c* is a decimal digit; otherwise, it shall return 0.

21146 **ERRORS**

21147            No errors are defined.

21148 **EXAMPLES**

21149            None.

21150 **APPLICATION USAGE**

21151            To ensure applications portability, especially across natural languages, only this function and  
21152            those listed in the SEE ALSO section should be used for character classification.

21153 **RATIONALE**

21154            None.

21155 **FUTURE DIRECTIONS**

21156            None.

21157 **SEE ALSO**

21158            *isalnum()*, *isalpha()*, *iscntrl()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
21159            *isxdigit()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>

21160 **CHANGE HISTORY**

21161            First released in Issue 1. Derived from Issue 1 of the SVID.

21162 **Issue 4**

21163            The text of the DESCRIPTION is revised, although there are no functional differences between  
21164            this issue and Issue 3.

21165 **Issue 6**

21166            The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21167 **NAME**

21168       isfdtype — determine whether a file descriptor refers to a socket

21169 **SYNOPSIS**

21170       #include <sys/stat.h>

21171       int isfdtype(int *fildev*, int *fdtype*);

21172 **DESCRIPTION**

21173       The *isfdtype()* function shall determine whether the descriptor *fildev* has the properties identified  
21174       by the value of *fdtype*, returning 1 if so, and 0 if not.

21175       If *fdtype* has the value S\_IFSOCK:

- 21176       • The *isfdtype()* function shall return 1 if the descriptor refers to a socket.
- 21177       • It is implementation-defined whether *isfdtype()* shall return 1 if the descriptor refers to a  
21178       pipe.
- 21179       • The function shall return 0 for descriptors that refer neither to a socket nor to a pipe.

21180 **RETURN VALUE**

21181       Upon successful completion, the *isfdtype()* function shall return a value of 1 or 0 indicating  
21182       whether the descriptor is of the indicated type. Otherwise, it shall return a value of -1 and set  
21183       *errno* to indicate the error.

21184 **ERRORS**

21185       If any of the following conditions occur, the *isfdtype()* function shall return -1 and set *errno* to  
21186       the corresponding value:

21187       [EBADF]       The *fildev* argument is not a valid file descriptor.

21188 **EXAMPLES**

21189       None.

21190 **APPLICATION USAGE**

21191       None.

21192 **RATIONALE**

21193       None.

21194 **FUTURE DIRECTIONS**

21195       None.

21196 **SEE ALSO**

21197       *isatty()*, *socket()*, *stat()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/stat.h>

21198 **CHANGE HISTORY**

21199       First released in Issue 6. Derived from IEEE Std. 1003.1g-2000.

21200 **NAME**

21201           isfinite — test for finite value

21202 **SYNOPSIS**

21203           #include <math.h>

21204           int isfinite(real-floating x);

21205 **DESCRIPTION**

21206 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21207       conflict between the requirements described here and the ISO C standard is unintentional. This  
21208       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21209       The *isfinite()* macro shall determine whether its argument has a finite value (zero, subnormal, or  
21210       normal, and not infinite or NaN). First, an argument represented in a format wider than its  
21211       semantic type is converted to its semantic type. Then determination is based on the type of the  
21212       argument.

21213 **RETURN VALUE**

21214       The *isfinite()* macro shall return a non-zero value if and only if its argument has a finite value.

21215 **ERRORS**

21216       No errors are defined.

21217 **EXAMPLES**

21218       None.

21219 **APPLICATION USAGE**

21220       None.

21221 **RATIONALE**

21222       None.

21223 **FUTURE DIRECTIONS**

21224       None.

21225 **SEE ALSO**

21226       *fpclassify()*, *isinf()*, *isnan()*, *isnormal()*, *signbit()*, the Base Definitions volume of  
21227       IEEE Std. 1003.1-200x <math.h>

21228 **CHANGE HISTORY**

21229       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21230 **NAME**

21231 isgraph — test for a visible character

21232 **SYNOPSIS**

21233 #include &lt;ctype.h&gt;

21234 int isgraph(int c);

21235 **DESCRIPTION**

21236 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
21237 conflict between the requirements described here and the ISO C standard is unintentional. This  
21238 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21239 The *isgraph()* function shall test whether *c* is a character of class **graph** in the program's current  
21240 locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

21241 In all cases *c* is an **int**, the value of which the application shall ensure is a character representable  
21242 as an **unsigned char** or equal to the value of the macro EOF. If the argument has any other value,  
21243 the behavior is undefined.

21244 **RETURN VALUE**

21245 The *isgraph()* function shall return non-zero if *c* is a character with a visible representation;  
21246 otherwise, it shall return 0.

21247 **ERRORS**

21248 No errors are defined.

21249 **EXAMPLES**

21250 None.

21251 **APPLICATION USAGE**

21252 To ensure applications portability, especially across natural languages, only this function and  
21253 those listed in the SEE ALSO section should be used for character classification.

21254 **RATIONALE**

21255 None.

21256 **FUTURE DIRECTIONS**

21257 None.

21258 **SEE ALSO**

21259 *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*,  
21260 *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base Definitions  
21261 volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

21262 **CHANGE HISTORY**

21263 First released in Issue 1. Derived from Issue 1 of the SVID.

21264 **Issue 4**

21265 The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
21266 functional differences between this issue and Issue 3. Operation in the C locale is no longer  
21267 described explicitly on this reference page.

21268 **Issue 6**

21269 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21270 **NAME**21271 `isgreater` — test if  $x$  greater than  $y$ 21272 **SYNOPSIS**21273 `#include <math.h>`21274 `int isgreater(real-floating x, real-floating y);`21275 **DESCRIPTION**

21276 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
21277 conflict between the requirements described here and the ISO C standard is unintentional. This  
21278 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21279 The `isgreater()` macro shall determine whether its first argument is greater than its second  
21280 argument. The value of `isgreater(x, y)` shall be equal to  $(x) > (y)$ ; however, unlike  $(x) > (y)$ ,  
21281 `isgreater(x, y)` shall not raise the invalid floating-point exception when  $x$  and  $y$  are unordered.

21282 **RETURN VALUE**21283 Upon successful completion, the `isgreater()` macro shall return the value of  $(x) > (y)$ .21284 If  $x$  or  $y$  is NaN, 0 shall be returned.21285 **ERRORS**

21286 No errors are defined.

21287 **EXAMPLES**

21288 None.

21289 **APPLICATION USAGE**

21290 The relational and equality operators support the usual mathematical relationships between  
21291 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21292 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21293 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21294 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21295 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21296 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21297 indicates that the argument shall be an expression of **real-floating** type.

21298 **RATIONALE**

21299 None.

21300 **FUTURE DIRECTIONS**

21301 None.

21302 **SEE ALSO**

21303 `isgreaterequal()`, `isless()`, `islessequal()`, `islessgreater()`, `isunordered()`, the Base Definitions volume of  
21304 IEEE Std. 1003.1-200x `<math.h>`

21305 **CHANGE HISTORY**

21306 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21307 **NAME**

21308           isgreaterqual — test if *x* greater than or equal to *y*

21309 **SYNOPSIS**

21310           #include <math.h>

21311           int isgreaterqual(real-floating *x*, real-floating *y*);

21312 **DESCRIPTION**

21313 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21314 conflict between the requirements described here and the ISO C standard is unintentional. This  
21315 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21316       The *isgreaterqual()* macro shall determine whether its first argument is greater than or equal to  
21317 its second argument. The value of *isgreaterqual(x, y)* shall be equal to  $(x) \geq (y)$ ; however, unlike  
21318  $(x) \geq (y)$ , *isgreaterqual(x, y)* shall not raise the invalid floating-point exception when *x* and *y* are  
21319 unordered.

21320 **RETURN VALUE**

21321       Upon successful completion, the *isgreaterqual()* macro shall return the value of  $(x) \geq (y)$ .

21322       If *x* or *y* is NaN, 0 shall be returned.

21323 **ERRORS**

21324       No errors are defined.

21325 **EXAMPLES**

21326       None.

21327 **APPLICATION USAGE**

21328       The relational and equality operators support the usual mathematical relationships between  
21329 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21330 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21331 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21332 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21333 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21334 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21335 indicates that the argument shall be an expression of **real-floating** type.

21336 **RATIONALE**

21337       None.

21338 **FUTURE DIRECTIONS**

21339       None.

21340 **SEE ALSO**

21341       *isgreater()*, *isless()*, *islessequal()*, *islessgreater()*, *isunordered()*, the Base Definitions volume of  
21342 IEEE Std. 1003.1-200x <math.h>

21343 **CHANGE HISTORY**

21344       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21345 **NAME**

21346           isinf — test for infinity

21347 **SYNOPSIS**

21348           #include &lt;math.h&gt;

21349           int isinf(real-floating x);

21350 **DESCRIPTION**

21351 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
21352       conflict between the requirements described here and the ISO C standard is unintentional. This  
21353       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21354       The *isinf()* macro shall determine whether its argument value is an infinity (positive or  
21355       negative). First, an argument represented in a format wider than its semantic type is converted  
21356       to its semantic type. Then determination is based on the type of the argument.

21357 **RETURN VALUE**21358       The *isinf()* macro shall return a non-zero value if and only if its argument has an infinite value.21359 **ERRORS**

21360       No errors are defined.

21361 **EXAMPLES**

21362       None.

21363 **APPLICATION USAGE**

21364       None.

21365 **RATIONALE**

21366       None.

21367 **FUTURE DIRECTIONS**

21368       None.

21369 **SEE ALSO**

21370       *fpclassify()*, *isfinite()*, *isnan()*, *isnormal()*, *signbit()*, the Base Definitions volume of  
21371       IEEE Std. 1003.1-200x <math.h>

21372 **CHANGE HISTORY**

21373       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21374 **NAME**

21375 isless — test if *x* is less than *y*

21376 **SYNOPSIS**

21377 #include <math.h>

21378 int isless(real-floating *x*, real-floating *y*);

21379 **DESCRIPTION**

21380 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
21381 conflict between the requirements described here and the ISO C standard is unintentional. This  
21382 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21383 The *isless()* macro shall determine whether its first argument is less than its second argument.  
21384 The value of *isless(x, y)* shall be equal to  $(x) < (y)$ ; however, unlike  $(x) < (y)$ , *isless(x, y)* shall not  
21385 raise the invalid floating-point exception when *x* and *y* are unordered.

21386 **RETURN VALUE**

21387 Upon successful completion, the *isless()* macro shall return the value of  $(x) < (y)$ .

21388 If *x* or *y* is NaN, 0 shall be returned.

21389 **ERRORS**

21390 No errors are defined.

21391 **EXAMPLES**

21392 None.

21393 **APPLICATION USAGE**

21394 The relational and equality operators support the usual mathematical relationships between  
21395 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21396 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21397 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21398 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21399 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21400 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21401 indicates that the argument shall be an expression of **real-floating** type.

21402 **RATIONALE**

21403 None.

21404 **FUTURE DIRECTIONS**

21405 None.

21406 **SEE ALSO**

21407 *isgreater()*, *isgreaterequal()*, *islessequal()*, *islessgreater()*, *isunordered()*, the Base Definitions volume  
21408 of IEEE Std. 1003.1-200x, <math.h>

21409 **CHANGE HISTORY**

21410 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21411 **NAME**

21412 islessequal — test if *x* is less than or equal to *y*

21413 **SYNOPSIS**

21414 #include <math.h>

21415 int islessequal(real-floating *x*, real-floating *y*);

21416 **DESCRIPTION**

21417 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
21418 conflict between the requirements described here and the ISO C standard is unintentional. This  
21419 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21420 The *islessequal()* macro shall determine whether its first argument is less than or equal to its  
21421 second argument. The value of *islessequal(x, y)* shall be equal to  $(x) \leq (y)$ ; however, unlike  
21422  $(x) \leq (y)$ , *islessequal(x, y)* shall not raise the invalid floating-point exception when *x* and *y* are  
21423 unordered.

21424 **RETURN VALUE**

21425 Upon successful completion, the *islessequal()* macro shall return the value of  $(x) \leq (y)$ .

21426 If *x* or *y* is NaN, 0 shall be returned.

21427 **ERRORS**

21428 No errors are defined.

21429 **EXAMPLES**

21430 None.

21431 **APPLICATION USAGE**

21432 The relational and equality operators support the usual mathematical relationships between  
21433 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21434 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21435 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21436 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21437 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21438 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21439 indicates that the argument shall be an expression of **real-floating** type.

21440 **RATIONALE**

21441 None.

21442 **FUTURE DIRECTIONS**

21443 None.

21444 **SEE ALSO**

21445 *isgreater()*, *isgreaterequal()*, *isless()*, *islessgreater()*, *isunordered()*, the Base Definitions volume of  
21446 IEEE Std. 1003.1-200x <math.h>

21447 **CHANGE HISTORY**

21448 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21449 **NAME**

21450 islessgreater — test if  $x$  is less than or greater than  $y$

21451 **SYNOPSIS**

21452 #include <math.h>

21453 int islessgreater(real-floating  $x$ , real-floating  $y$ );

21454 **DESCRIPTION**

21455 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
21456 conflict between the requirements described here and the ISO C standard is unintentional. This  
21457 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21458 The *islessgreater()* macro shall determine whether its first argument is less than or greater than  
21459 its second argument. The *islessgreater*( $x$ ,  $y$ ) macro is similar to  $(x) < (y) \ || \ (x) > (y)$ ; however,  
21460 *islessgreater*( $x$ ,  $y$ ) shall not raise the invalid floating-point exception when  $x$  and  $y$  are unordered  
21461 (nor shall it evaluate  $x$  and  $y$  twice).

21462 **RETURN VALUE**

21463 Upon successful completion, the *islessgreater()* macro shall return the value of  
21464  $(x) < (y) \ || \ (x) > (y)$ .

21465 If  $x$  or  $y$  is NaN, 0 shall be returned.

21466 **ERRORS**

21467 No errors are defined.

21468 **EXAMPLES**

21469 None.

21470 **APPLICATION USAGE**

21471 The relational and equality operators support the usual mathematical relationships between  
21472 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21473 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21474 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21475 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21476 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21477 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21478 indicates that the argument shall be an expression of **real-floating** type.

21479 **RATIONALE**

21480 None.

21481 **FUTURE DIRECTIONS**

21482 None.

21483 **SEE ALSO**

21484 *isgreater()*, *isgreaterequal()*, *isless()*, *islessequal()*, *isunordered()*, the Base Definitions volume of  
21485 IEEE Std. 1003.1-200x <math.h>

21486 **CHANGE HISTORY**

21487 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21488 **NAME**

21489           islower — test for a lowercase letter

21490 **SYNOPSIS**

21491           #include &lt;ctype.h&gt;

21492           int islower(int c);

21493 **DESCRIPTION**

21494 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 21495 conflict between the requirements described here and the ISO C standard is unintentional. This  
 21496 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21497       The *islower()* function shall test whether *c* is a character of class **lower** in the program's current  
 21498 locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

21499       In all cases *c* is an **int**, the value of which the application shall ensure is a character representable  
 21500 as an **unsigned char** or equal to the value of the macro EOF. If the argument has any other value,  
 21501 the behavior is undefined.

21502 **RETURN VALUE**21503       The *islower()* function shall return non-zero if *c* is a lowercase letter; otherwise, it shall return 0.21504 **ERRORS**

21505       No errors are defined.

21506 **EXAMPLES**21507           **Testing for a Lowercase Letter**

21508       The following example tests whether the value is a lowercase letter, based on the locale of the  
 21509 user, then uses it as part of a key value.

```

21510 #include <ctype.h>
21511 #include <stdlib.h>
21512 #include <locale.h>
21513 ...
21514 char *keyst;
21515 int elementlen, len;
21516 char c;
21517 ...
21518 setlocale(LC_ALL, "");
21519 ...
21520 len = 0;
21521 while (len < elementlen) {
21522 c = (char) (rand() % 256);
21523 ...
21524 if (islower(c))
21525 keyst[len++] = c;
21526 }
21527 ...
```

21528 **APPLICATION USAGE**

21529       To ensure applications portability, especially across natural languages, only this function and  
 21530 those listed in the SEE ALSO section should be used for character classification.

21531 **RATIONALE**

21532           None.

21533 **FUTURE DIRECTIONS**

21534           None.

21535 **SEE ALSO**21536           *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,21537           *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base

21538           Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

21539 **CHANGE HISTORY**

21540           First released in Issue 1. Derived from Issue 1 of the SVID.

21541 **Issue 4**21542           The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
21543           functional differences between this issue and Issue 3. Operation in the C locale is no longer  
21544           described explicitly on this reference page.21545 **Issue 6**

21546           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21547 **NAME**21548        **isnan** — test for a NaN21549 **SYNOPSIS**

21550        #include &lt;math.h&gt;

21551        int isnan(real-floating x);

21552 **DESCRIPTION**

21553 **cx**        The functionality described on this reference page is aligned with the ISO C standard. Any  
21554        conflict between the requirements described here and the ISO C standard is unintentional. This  
21555        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21556        The *isnan()* macro shall determine whether its argument value is a NaN. First, an argument  
21557        represented in a format wider than its semantic type is converted to its semantic type. Then  
21558        determination is based on the type of the argument.

21559 **RETURN VALUE**21560        The *isnan()* macro shall return a non-zero value if and only if its argument has a NaN value.21561 **ERRORS**

21562        No errors are defined.

21563 **EXAMPLES**

21564        None.

21565 **APPLICATION USAGE**

21566        None.

21567 **RATIONALE**

21568        None.

21569 **FUTURE DIRECTIONS**

21570        None.

21571 **SEE ALSO**

21572        *fpclassify()*, *isfinite()*, *isinf()*, *isnormal()*, *signbit()*, the Base Definitions volume of  
21573        IEEE Std. 1003.1-200x, <math.h>

21574 **CHANGE HISTORY**

21575        First released in Issue 3.

21576 **Issue 4**

21577        The words “not supporting NaN” are added to the APPLICATION USAGE section.

21578 **Issue 5**

21579        The DESCRIPTION is updated to indicate the return value when NaN is not supported. This  
21580        text was previously published in the APPLICATION USAGE section.

21581 **Issue 6**

21582        Entry re-written for alignment with the ISO/IEC 9899:1999 standard.

21583 **NAME**

21584           isnormal — test for a normal value

21585 **SYNOPSIS**

21586           #include <math.h>

21587           int isnormal(real-floating x);

21588 **DESCRIPTION**

21589 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21590       conflict between the requirements described here and the ISO C standard is unintentional. This  
21591       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21592       The *isnormal()* macro shall determine whether its argument value is normal (neither zero,  
21593       subnormal, infinite, nor NaN). First, an argument represented in a format wider than its  
21594       semantic type is converted to its semantic type. Then determination is based on the type of the  
21595       argument.

21596 **RETURN VALUE**

21597       The *isnormal()* macro shall return a non-zero value if and only if its argument has a normal  
21598       value.

21599 **ERRORS**

21600       No errors are defined.

21601 **EXAMPLES**

21602       None.

21603 **APPLICATION USAGE**

21604       None.

21605 **RATIONALE**

21606       None.

21607 **FUTURE DIRECTIONS**

21608       None.

21609 **SEE ALSO**

21610       *fpclassify()*, *isfinite()*, *isinf()*, *isnan()*, *signbit()*, the Base Definitions volume of  
21611       IEEE Std. 1003.1-200x, <math.h>

21612 **CHANGE HISTORY**

21613       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21614 **NAME**

21615           isprint — test for a printing character

21616 **SYNOPSIS**

21617           #include <ctype.h>

21618           int isprint(int c);

21619 **DESCRIPTION**

21620 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21621       conflict between the requirements described here and the ISO C standard is unintentional. This  
21622       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21623       The *isprint()* function shall test whether *c* is a character of class **print** in the program's current  
21624       locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

21625       In all cases *c* is an **int**, the value of which the application shall ensure is a character representable  
21626       as an **unsigned char** or equal to the value of the macro EOF. If the argument has any other value,  
21627       the behavior is undefined.

21628 **RETURN VALUE**

21629       The *isprint()* function shall return non-zero if *c* is a printing character; otherwise, it shall return 0.

21630 **ERRORS**

21631       No errors are defined.

21632 **EXAMPLES**

21633       None.

21634 **APPLICATION USAGE**

21635       To ensure applications portability, especially across natural languages, only this function and  
21636       those listed in the SEE ALSO section should be used for character classification.

21637 **RATIONALE**

21638       None.

21639 **FUTURE DIRECTIONS**

21640       None.

21641 **SEE ALSO**

21642       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *ispunct()*, *isspace()*, *isupper()*,  
21643       *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base  
21644       Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

21645 **CHANGE HISTORY**

21646       First released in Issue 1. Derived from Issue 1 of the SVID.

21647 **Issue 4**

21648       The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
21649       functional differences between this issue and Issue 3. Operation in the C locale is no longer  
21650       described explicitly on this reference page.

21651 **Issue 6**

21652       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21653 **NAME**

21654           ispunct — test for a punctuation character

21655 **SYNOPSIS**

21656           #include &lt;ctype.h&gt;

21657           int ispunct(int c);

21658 **DESCRIPTION**

21659 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21660 conflict between the requirements described here and the ISO C standard is unintentional. This  
21661 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21662       The *ispunct()* function shall test whether *c* is a character of class **punct** in the program's current  
21663 locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

21664       In all cases *c* is an **int**, the value of which the application shall ensure is a character representable  
21665 as an **unsigned char** or equal to the value of the macro EOF. If the argument has any other value,  
21666 the behavior is undefined.

21667 **RETURN VALUE**

21668       The *ispunct()* function shall return non-zero if *c* is a punctuation character; otherwise, it shall  
21669 return 0.

21670 **ERRORS**

21671       No errors are defined.

21672 **EXAMPLES**

21673       None.

21674 **APPLICATION USAGE**

21675       To ensure applications portability, especially across natural languages, only this function and  
21676 those listed in the SEE ALSO section should be used for character classification.

21677 **RATIONALE**

21678       None.

21679 **FUTURE DIRECTIONS**

21680       None.

21681 **SEE ALSO**

21682       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *isspace()*, *isupper()*, *isxdigit()*,  
21683 *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base Definitions  
21684 volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

21685 **CHANGE HISTORY**

21686       First released in Issue 1. Derived from Issue 1 of the SVID.

21687 **Issue 4**

21688       The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
21689 functional differences between this issue and Issue 3. Operation in the C locale is no longer  
21690 described explicitly on this reference page.

21691 **Issue 6**

21692       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21693 **NAME**

21694           isspace — test for a white-space character

21695 **SYNOPSIS**

21696           #include <ctype.h>

21697           int isspace(int c);

21698 **DESCRIPTION**

21699 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21700 conflict between the requirements described here and the ISO C standard is unintentional. This  
21701 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21702       The *isspace()* function shall test whether *c* is a character of class **space** in the program's current  
21703 locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

21704       In all cases *c* is an **int**, the value of which the application shall ensure is a character representable  
21705 as an **unsigned char** or equal to the value of the macro EOF. If the argument has any other value,  
21706 the behavior is undefined.

21707 **RETURN VALUE**

21708       The *isspace()* function shall return non-zero if *c* is a white-space character; otherwise, it shall  
21709 return 0.

21710 **ERRORS**

21711       No errors are defined.

21712 **EXAMPLES**

21713       None.

21714 **APPLICATION USAGE**

21715       To ensure applications portability, especially across natural languages, only this function and  
21716 those listed in the SEE ALSO section should be used for character classification.

21717 **RATIONALE**

21718       None.

21719 **FUTURE DIRECTIONS**

21720       None.

21721 **SEE ALSO**

21722       *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isupper()*,  
21723 *isxdigit()*, *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base  
21724 Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

21725 **CHANGE HISTORY**

21726       First released in Issue 1. Derived from Issue 1 of the SVID.

21727 **Issue 4**

21728       The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
21729 functional differences between this issue and Issue 3. Operation in the C locale is no longer  
21730 described explicitly on this reference page.

21731 **Issue 6**

21732       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21733 **NAME**

21734 isunordered — test if arguments are unordered

21735 **SYNOPSIS**

21736 #include <math.h>

21737 int isunordered(real-floating x, real-floating y);

21738 **DESCRIPTION**

21739 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
21740 conflict between the requirements described here and the ISO C standard is unintentional. This  
21741 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21742 The *isunordered()* macro shall determine whether its arguments are unordered.

21743 **RETURN VALUE**

21744 Upon successful completion, the *isunordered()* macro shall return 1 if its arguments are  
21745 unordered, and 0 otherwise.

21746 If *x* or *y* is NaN, 0 shall be returned.

21747 **ERRORS**

21748 No errors are defined.

21749 **EXAMPLES**

21750 None.

21751 **APPLICATION USAGE**

21752 The relational and equality operators support the usual mathematical relationships between  
21753 numeric values. For any ordered pair of numeric values, exactly one of the relationships (less,  
21754 greater, and equal) is true. Relational operators may raise the invalid floating-point exception  
21755 when argument values are NaNs. For a NaN and a numeric value, or for two NaNs, just the  
21756 unordered relationship is true. This macro is a quiet (non-floating-point exception raising)  
21757 version of a relational operator. It facilitates writing efficient code that accounts for NaNs  
21758 without suffering the invalid floating-point exception. In the SYNOPSIS section, **real-floating**  
21759 indicates that the argument shall be an expression of **real-floating** type.

21760 **RATIONALE**

21761 None.

21762 **FUTURE DIRECTIONS**

21763 None.

21764 **SEE ALSO**

21765 *isgreater()*, *isgreaterequal()*, *isless()*, *islessequal()*, *islessgreater()*, the Base Definitions volume of  
21766 IEEE Std. 1003.1-200x, <math.h>

21767 **CHANGE HISTORY**

21768 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21769 **NAME**

21770 isupper — test for an uppercase letter

21771 **SYNOPSIS**

21772 #include <ctype.h>

21773 int isupper(int c);

21774 **DESCRIPTION**

21775 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
21776 conflict between the requirements described here and the ISO C standard is unintentional. This  
21777 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21778 The *isupper()* function shall test whether *c* is a character of class **upper** in the program's current  
21779 locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

21780 In all cases *c* is an **int**, the value of which the application shall ensure is a character representable  
21781 as an **unsigned char** or equal to the value of the macro EOF. If the argument has any other value,  
21782 the behavior is undefined.

21783 **RETURN VALUE**

21784 The *isupper()* function shall return non-zero if *c* is an uppercase letter; otherwise, it shall return 0.

21785 **ERRORS**

21786 No errors are defined.

21787 **EXAMPLES**

21788 None.

21789 **APPLICATION USAGE**

21790 To ensure applications portability, especially across natural languages, only this function and  
21791 those listed in the SEE ALSO section should be used for character classification.

21792 **RATIONALE**

21793 None.

21794 **FUTURE DIRECTIONS**

21795 None.

21796 **SEE ALSO**

21797 *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isxdigit()*,  
21798 *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base Definitions  
21799 volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

21800 **CHANGE HISTORY**

21801 First released in Issue 1. Derived from Issue 1 of the SVID.

21802 **Issue 4**

21803 The text of the DESCRIPTION and RETURN VALUE sections is revised, although there are no  
21804 functional differences between this issue and Issue 3. Operation in the C locale is no longer  
21805 described explicitly on this reference page.

21806 **Issue 6**

21807 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21808 **NAME**

21809           iswalnum — test for an alphanumeric wide-character code

21810 **SYNOPSIS**

21811           #include <wctype.h>

21812           int iswalnum(wint\_t *wc*);

21813 **DESCRIPTION**

21814 *cx*       The functionality described on this reference page is aligned with the ISO C standard. Any  
21815 conflict between the requirements described here and the ISO C standard is unintentional. This  
21816 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21817       The *iswalnum()* function shall test whether *wc* is a wide-character code representing a character  
21818 of class **alpha** or **digit** in the program's current locale; see the Base Definitions volume of  
21819 IEEE Std. 1003.1-200x, Chapter 7, Locale.

21820       In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21821 code corresponding to a valid character in the current locale, or equal to the value of the macro  
21822 WEOF. If the argument has any other value, the behavior is undefined.

21823 **RETURN VALUE**

21824       The *iswalnum()* function shall return non-zero if *wc* is an alphanumeric wide-character code;  
21825 otherwise, it shall return 0.

21826 **ERRORS**

21827       No errors are defined.

21828 **EXAMPLES**

21829       None.

21830 **APPLICATION USAGE**

21831       To ensure applications portability, especially across natural languages, only this function and  
21832 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21833 **RATIONALE**

21834       None.

21835 **FUTURE DIRECTIONS**

21836       None.

21837 **SEE ALSO**

21838       *iswalphabet()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
21839 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21840 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>, <stdio.h>, the Base Definitions volume of  
21841 IEEE Std. 1003.1-200x, Chapter 7, Locale

21842 **CHANGE HISTORY**

21843       First released as a World-wide Portability Interface in Issue 4.

21844 **Issue 5**

21845       The following change has been made in this issue for alignment with  
21846 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21847       • The SYNOPSIS has been changed to indicate that this function and associated data types are  
21848       now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

21849 **Issue 6**

21850

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21851 **NAME**

21852           iswalpha — test for an alphabetic wide-character code

21853 **SYNOPSIS**

21854           #include <wctype.h>

21855           int iswalpha(wint\_t *wc*);

21856 **DESCRIPTION**

21857 *cx*       The functionality described on this reference page is aligned with the ISO C standard. Any  
21858 conflict between the requirements described here and the ISO C standard is unintentional. This  
21859 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21860       The *iswalpha()* function shall test whether *wc* is a wide-character code representing a character of  
21861 class **alpha** in the program's current locale; see the Base Definitions volume of  
21862 IEEE Std. 1003.1-200x, Chapter 7, Locale.

21863       In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21864 code corresponding to a valid character in the current locale, or equal to the value of the macro  
21865 WEOF. If the argument has any other value, the behavior is undefined.

21866 **RETURN VALUE**

21867       The *iswalpha()* function shall return non-zero if *wc* is an alphabetic wide-character code;  
21868 otherwise, it shall return 0.

21869 **ERRORS**

21870       No errors are defined.

21871 **EXAMPLES**

21872       None.

21873 **APPLICATION USAGE**

21874       To ensure applications portability, especially across natural languages, only this function and  
21875 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21876 **RATIONALE**

21877       None.

21878 **FUTURE DIRECTIONS**

21879       None.

21880 **SEE ALSO**

21881       *iswalnum()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
21882 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21883 IEEE Std. 1003.1-200x, <**wctype.h**>, <**wchar.h**>, <**stdio.h**>, the Base Definitions volume of  
21884 IEEE Std. 1003.1-200x, Chapter 7, Locale

21885 **CHANGE HISTORY**

21886       First released in Issue 4.

21887 **Issue 5**

21888       The following change has been made in this issue for alignment with  
21889 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21890       • The SYNOPSIS has been changed to indicate that this function and associated data types are  
21891       now made visible by inclusion of the <**wctype.h**> header rather than <**wchar.h**>.

21892 **Issue 6**

21893

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21894 **NAME**

21895           iswblank — test for a blank wide-character code

21896 **SYNOPSIS**

21897           #include <wctype.h>

21898           int iswblank(wint\_t *wc*);

21899 **DESCRIPTION**

21900 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
21901 conflict between the requirements described here and the ISO C standard is unintentional. This  
21902 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21903       The *iswblank()* function shall test whether *wc* is a wide-character code representing a character of  
21904 class **blank** in the program's current locale; see the Base Definitions volume of  
21905 IEEE Std. 1003.1-200x, Chapter 7, Locale. In all cases, *wc* is a **wint\_t**, the value of which the  
21906 application shall ensure is a wide-character code corresponding to a valid character in the  
21907 current locale, or equal to the value of the macro WEOF. If the argument has any other value, the  
21908 behavior is undefined.

21909 **RETURN VALUE**

21910       The *iswblank()* function shall return non-zero if *wc* is a <blank> wide-character code; otherwise,  
21911 it shall return 0.

21912 **ERRORS**

21913       No errors are defined.

21914 **EXAMPLES**

21915       None.

21916 **APPLICATION USAGE**

21917       To ensure applications portability, especially across natural languages, only this function and  
21918 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21919 **RATIONALE**

21920       None.

21921 **FUTURE DIRECTIONS**

21922       None.

21923 **SEE ALSO**

21924       *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
21925 *iswpunct()*, *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
21926 IEEE Std. 1003.1-200x, <wchar.h>, <wctype.h>, <stdio.h>

21927 **CHANGE HISTORY**

21928       First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

21929 **NAME**

21930 `iswcntrl` — test for a control wide-character code

21931 **SYNOPSIS**

21932 `#include <wctype.h>`

21933 `int iswcntrl(wint_t wc);`

21934 **DESCRIPTION**

21935 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
21936 conflict between the requirements described here and the ISO C standard is unintentional. This  
21937 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21938 The `iswcntrl()` function shall test whether `wc` is a wide-character code representing a character of  
21939 class **control** in the program's current locale; see the Base Definitions volume of  
21940 IEEE Std. 1003.1-200x, Chapter 7, Locale.

21941 In all cases `wc` is a **wint\_t**, the value of which the application shall ensure is a wide-character  
21942 code corresponding to a valid character in the current locale, or equal to the value of the macro  
21943 `WEOF`. If the argument has any other value, the behavior is undefined.

21944 **RETURN VALUE**

21945 The `iswcntrl()` function shall return non-zero if `wc` is a control wide-character code; otherwise, it  
21946 shall return 0.

21947 **ERRORS**

21948 No errors are defined.

21949 **EXAMPLES**

21950 None.

21951 **APPLICATION USAGE**

21952 To ensure applications portability, especially across natural languages, only this function and  
21953 those listed in the SEE ALSO section should be used for classification of wide-character codes.

21954 **RATIONALE**

21955 None.

21956 **FUTURE DIRECTIONS**

21957 None.

21958 **SEE ALSO**

21959 `iswalnum()`, `iswalpha()`, `iswctype()`, `iswdigit()`, `iswgraph()`, `iswlower()`, `iswprint()`, `iswpunct()`,  
21960 `iswspace()`, `iswupper()`, `iswxdigit()`, `setlocale()`, the Base Definitions volume of  
21961 IEEE Std. 1003.1-200x, `<wctype.h>`, `<wchar.h>`, the Base Definitions volume of  
21962 IEEE Std. 1003.1-200x, Chapter 7, Locale

21963 **CHANGE HISTORY**

21964 First released in Issue 4.

21965 **Issue 5**

21966 The following change has been made in this issue for alignment with  
21967 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 21968 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
21969 now made visible by inclusion of the `<wctype.h>` header rather than `<wchar.h>`.

21970 **Issue 6**

21971

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

21972 **NAME**

21973 iswctype — test character for a specified class

21974 **SYNOPSIS**

21975 #include &lt;wctype.h&gt;

21976 int iswctype(wint\_t *wc*, wctype\_t *charclass*);21977 **DESCRIPTION**

21978 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 21979 conflict between the requirements described here and the ISO C standard is unintentional. This  
 21980 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

21981 The *iswctype()* function shall determine whether the wide-character code *wc* has the character  
 21982 class *charclass*, returning true or false. The *iswctype()* function is defined on WEOF and wide-  
 21983 character codes corresponding to the valid character encodings in the current locale. If the *wc*  
 21984 argument is not in the domain of the function, the result is undefined. If the value of *charclass* is  
 21985 invalid (that is, not obtained by a call to *wctype()* or *charclass* is invalidated by a subsequent call  
 21986 to *setlocale()* that has affected category *LC\_CTYPE*) the result is unspecified.

21987 **RETURN VALUE**

21988 The *iswctype()* function shall return non-zero (true) if and only if *wc* has the property described  
 21989 **CX** by *charclass*. If *charclass* is 0, *iswctype()* shall return 0.

21990 **ERRORS**

21991 No errors are defined.

21992 **EXAMPLES**21993 **Testing for a Valid Character**

```
21994 #include <wctype.h>
21995 ...
21996 int yes_or_no;
21997 wint_t wc;
21998 wctype_t valid_class;
21999 ...
22000 if ((valid_class=wctype("vowel")) == (wctype_t)0)
22001 /* Invalid character class. */
22002 yes_or_no=iswctype(wc,valid_class);
```

22003 **APPLICATION USAGE**

22004 The twelve strings "alnum", "alpha", "blank", "cntrl", "digit", "graph", "lower",  
 22005 "print", "punct", "space", "upper", and "xdigit" are reserved for the standard  
 22006 character classes. In the table below, the functions in the left column are equivalent to the  
 22007 functions in the right column.

|       |                     |                                      |
|-------|---------------------|--------------------------------------|
| 22008 | <i>iswalnum(wc)</i> | <i>iswctype(wc, wctype("alnum"))</i> |
| 22009 | <i>iswalpha(wc)</i> | <i>iswctype(wc, wctype("alpha"))</i> |
| 22010 | <i>iswblank(wc)</i> | <i>iswctype(wc, wctype("blank"))</i> |
| 22011 | <i>iswcntrl(wc)</i> | <i>iswctype(wc, wctype("cntrl"))</i> |
| 22012 | <i>iswdigit(wc)</i> | <i>iswctype(wc, wctype("digit"))</i> |
| 22013 | <i>iswgraph(wc)</i> | <i>iswctype(wc, wctype("graph"))</i> |
| 22014 | <i>iswlower(wc)</i> | <i>iswctype(wc, wctype("lower"))</i> |
| 22015 | <i>iswprint(wc)</i> | <i>iswctype(wc, wctype("print"))</i> |
| 22016 | <i>iswpunct(wc)</i> | <i>iswctype(wc, wctype("punct"))</i> |
| 22017 | <i>iswspace(wc)</i> | <i>iswctype(wc, wctype("space"))</i> |

22018            `iswupper(wc)`      `iswctype(wc, wctype("upper"))`  
22019            `iswxdigit(wc)`      `iswctype(wc, wctype("xdigit"))`

22020 **RATIONALE**

22021            None.

22022 **FUTURE DIRECTIONS**

22023            None.

22024 **SEE ALSO**

22025            `iswalnum()`, `iswalpha()`, `iswcntrl()`, `iswdigit()`, `iswgraph()`, `iswlower()`, `iswprint()`, `iswpunct()`,  
22026            `iswspace()`, `iswupper()`, `iswxdigit()`, `setlocale()`, `wctype()`, the Base Definitions volume of  
22027            IEEE Std. 1003.1-200x, `<wctype.h>`, `<wchar.h>`

22028 **CHANGE HISTORY**

22029            First released as World-wide Portability Interfaces in Issue 4.

22030 **Issue 5**

22031            The following change has been made in this issue for alignment with  
22032            ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22033            • The SYNOPSIS has been changed to indicate that this function and associated data types are  
22034            now made visible by inclusion of the `<wctype.h>` header rather than `<wchar.h>`.

22035 **NAME**

22036 iswdigit — test for a decimal digit wide-character code

22037 **SYNOPSIS**

22038 #include &lt;wctype.h&gt;

22039 int iswdigit(wint\_t *wc*);22040 **DESCRIPTION**

22041 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
 22042 conflict between the requirements described here and the ISO C standard is unintentional. This  
 22043 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22044 The *iswdigit()* function shall test whether *wc* is a wide-character code representing a character of  
 22045 class **digit** in the program's current locale; see the Base Definitions volume of  
 22046 IEEE Std. 1003.1-200x, Chapter 7, Locale.

22047 In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
 22048 code corresponding to a valid character in the current locale, or equal to the value of the macro  
 22049 WEOF. If the argument has any other value, the behavior is undefined.

22050 **RETURN VALUE**

22051 The *iswdigit()* function shall return non-zero if *wc* is a decimal digit wide-character code;  
 22052 otherwise, it shall return 0.

22053 **ERRORS**

22054 No errors are defined.

22055 **EXAMPLES**

22056 None.

22057 **APPLICATION USAGE**

22058 To ensure applications portability, especially across natural languages, only this function and  
 22059 those listed in the SEE ALSO section should be used for classification of wide-character codes.

22060 **RATIONALE**

22061 None.

22062 **FUTURE DIRECTIONS**

22063 None.

22064 **SEE ALSO**

22065 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
 22066 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
 22067 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>

22068 **CHANGE HISTORY**

22069 First released in Issue 4.

22070 **Issue 5**

22071 The following change has been made in this issue for alignment with  
 22072 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22073 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
 22074 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

22075 **Issue 6**

22076 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22077 **NAME**

22078 iswgraph — test for a visible wide-character code

22079 **SYNOPSIS**

22080 #include <wctype.h>

22081 int iswgraph(wint\_t wc);

22082 **DESCRIPTION**

22083 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
22084 conflict between the requirements described here and the ISO C standard is unintentional. This  
22085 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22086 The *iswgraph()* function shall test whether *wc* is a wide-character code representing a character  
22087 of class **graph** in the program's current locale; see the Base Definitions volume of  
22088 IEEE Std. 1003.1-200x, Chapter 7, Locale.

22089 In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
22090 code corresponding to a valid character in the current locale, or equal to the value of the macro  
22091 WEOF. If the argument has any other value, the behavior is undefined.

22092 **RETURN VALUE**

22093 The *iswgraph()* function shall return non-zero if *wc* is a wide-character code with a visible  
22094 representation; otherwise, it shall return 0.

22095 **ERRORS**

22096 No errors are defined.

22097 **EXAMPLES**

22098 None.

22099 **APPLICATION USAGE**

22100 To ensure applications portability, especially across natural languages, only this function and  
22101 those listed in the SEE ALSO section should be used for classification of wide-character codes.

22102 **RATIONALE**

22103 None.

22104 **FUTURE DIRECTIONS**

22105 None.

22106 **SEE ALSO**

22107 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
22108 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
22109 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
22110 IEEE Std. 1003.1-200x, Chapter 7, Locale

22111 **CHANGE HISTORY**

22112 First released in Issue 4.

22113 **Issue 5**

22114 The following change has been made in this issue for alignment with  
22115 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22116 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
22117 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

22118 **Issue 6**

22119

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22120 **NAME**

22121 iswlower — test for a lowercase letter wide-character code

22122 **SYNOPSIS**

22123 #include <wctype.h>

22124 int iswlower(wint\_t wc);

22125 **DESCRIPTION**

22126 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
22127 conflict between the requirements described here and the ISO C standard is unintentional. This  
22128 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22129 The *iswlower()* function shall test whether *wc* is a wide-character code representing a character  
22130 of class **lower** in the program's current locale; see the Base Definitions volume of  
22131 IEEE Std. 1003.1-200x, Chapter 7, Locale.

22132 In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
22133 code corresponding to a valid character in the current locale, or equal to the value of the macro  
22134 WEOF. If the argument has any other value, the behavior is undefined.

22135 **RETURN VALUE**

22136 The *iswlower()* function shall return non-zero if *wc* is a lowercase letter wide-character code;  
22137 otherwise, it shall return 0.

22138 **ERRORS**

22139 No errors are defined.

22140 **EXAMPLES**

22141 None.

22142 **APPLICATION USAGE**

22143 To ensure applications portability, especially across natural languages, only this function and  
22144 those listed in the SEE ALSO section should be used for classification of wide-character codes.

22145 **RATIONALE**

22146 None.

22147 **FUTURE DIRECTIONS**

22148 None.

22149 **SEE ALSO**

22150 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswprint()*, *iswpunct()*,  
22151 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
22152 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
22153 IEEE Std. 1003.1-200x, Chapter 7, Locale

22154 **CHANGE HISTORY**

22155 First released in Issue 4.

22156 **Issue 5**

22157 The following change has been made in this issue for alignment with  
22158 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22159 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
22160 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

22161 **Issue 6**

22162

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22163 **NAME**

22164 iswprint — test for a printing wide-character code

22165 **SYNOPSIS**

22166 #include &lt;wctype.h&gt;

22167 int iswprint(wint\_t *wc*);22168 **DESCRIPTION**

22169 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
22170 conflict between the requirements described here and the ISO C standard is unintentional. This  
22171 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22172 The *iswprint()* function shall test whether *wc* is a wide-character code representing a character of  
22173 class **print** in the program's current locale; see the Base Definitions volume of  
22174 IEEE Std. 1003.1-200x, Chapter 7, Locale.

22175 In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
22176 code corresponding to a valid character in the current locale, or equal to the value of the macro  
22177 WEOF. If the argument has any other value, the behavior is undefined.

22178 **RETURN VALUE**

22179 The *iswprint()* function shall return non-zero if *wc* is a printing wide-character code; otherwise, it  
22180 shall return 0.

22181 **ERRORS**

22182 No errors are defined.

22183 **EXAMPLES**

22184 None.

22185 **APPLICATION USAGE**

22186 To ensure applications portability, especially across natural languages, only this function and  
22187 those listed in the SEE ALSO section should be used for classification of wide-character codes.

22188 **RATIONALE**

22189 None.

22190 **FUTURE DIRECTIONS**

22191 None.

22192 **SEE ALSO**

22193 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswpunct()*,  
22194 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
22195 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
22196 IEEE Std. 1003.1-200x, Chapter 7, Locale

22197 **CHANGE HISTORY**

22198 First released in Issue 4.

22199 **Issue 5**

22200 The following change has been made in this issue for alignment with  
22201 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22202 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
22203 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

22204 **Issue 6**

22205

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22206 **NAME**

22207 iswpunct — test for a punctuation wide-character code

22208 **SYNOPSIS**

22209 #include <wctype.h>

22210 int iswpunct(wint\_t *wc*);

22211 **DESCRIPTION**

22212 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
22213 conflict between the requirements described here and the ISO C standard is unintentional. This  
22214 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22215 The *iswpunct()* function shall test whether *wc* is a wide-character code representing a character  
22216 of class **punct** in the program's current locale; see the Base Definitions volume of  
22217 IEEE Std. 1003.1-200x, Chapter 7, Locale.

22218 In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
22219 code corresponding to a valid character in the current locale, or equal to the value of the macro  
22220 WEOF. If the argument has any other value, the behavior is undefined.

22221 **RETURN VALUE**

22222 The *iswpunct()* function shall return non-zero if *wc* is a punctuation wide-character code;  
22223 otherwise, it shall return 0.

22224 **ERRORS**

22225 No errors are defined.

22226 **EXAMPLES**

22227 None.

22228 **APPLICATION USAGE**

22229 To ensure applications portability, especially across natural languages, only this function and  
22230 those listed in the SEE ALSO section should be used for classification of wide-character codes.

22231 **RATIONALE**

22232 None.

22233 **FUTURE DIRECTIONS**

22234 None.

22235 **SEE ALSO**

22236 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
22237 *iswspace()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
22238 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
22239 IEEE Std. 1003.1-200x, Chapter 7, Locale

22240 **CHANGE HISTORY**

22241 First released in Issue 4.

22242 **Issue 5**

22243 The following change has been made in this issue for alignment with  
22244 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22245 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
22246 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

22247 **Issue 6**

22248

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22249 **NAME**

22250 iswspace — test for a white-space wide-character code

22251 **SYNOPSIS**

22252 #include <wctype.h>

22253 int iswspace(wint\_t wc);

22254 **DESCRIPTION**

22255 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
22256 conflict between the requirements described here and the ISO C standard is unintentional. This  
22257 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22258 The *iswspace()* function shall test whether *wc* is a wide-character code representing a character of  
22259 class **space** in the program's current locale; see the Base Definitions volume of  
22260 IEEE Std. 1003.1-200x, Chapter 7, Locale.

22261 In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
22262 code corresponding to a valid character in the current locale, or equal to the value of the macro  
22263 WEOF. If the argument has any other value, the behavior is undefined.

22264 **RETURN VALUE**

22265 The *iswspace()* function shall return non-zero if *wc* is a white-space wide-character code;  
22266 otherwise, it shall return 0.

22267 **ERRORS**

22268 No errors are defined.

22269 **EXAMPLES**

22270 None.

22271 **APPLICATION USAGE**

22272 To ensure applications portability, especially across natural languages, only this function and  
22273 those listed in the SEE ALSO section should be used for classification of wide-character codes.

22274 **RATIONALE**

22275 None.

22276 **FUTURE DIRECTIONS**

22277 None.

22278 **SEE ALSO**

22279 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
22280 *iswpunct()*, *iswupper()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
22281 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
22282 IEEE Std. 1003.1-200x, Chapter 7, Locale

22283 **CHANGE HISTORY**

22284 First released in Issue 4.

22285 **Issue 5**

22286 The following change has been made in this issue for alignment with  
22287 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22288 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
22289 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

22290 **Issue 6**

22291

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22292 **NAME**

22293 iswupper — test for an uppercase letter wide-character code

22294 **SYNOPSIS**

22295 #include <wctype.h>

22296 int iswupper(wint\_t wc);

22297 **DESCRIPTION**

22298 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
22299 conflict between the requirements described here and the ISO C standard is unintentional. This  
22300 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22301 The *iswupper()* function shall test whether *wc* is a wide-character code representing a character  
22302 of class **upper** in the program's current locale; see the Base Definitions volume of  
22303 IEEE Std. 1003.1-200x, Chapter 7, Locale.

22304 In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
22305 code corresponding to a valid character in the current locale, or equal to the value of the macro  
22306 WEOF. If the argument has any other value, the behavior is undefined.

22307 **RETURN VALUE**

22308 The *iswupper()* function shall return non-zero if *wc* is an uppercase letter wide-character code;  
22309 otherwise, it shall return 0.

22310 **ERRORS**

22311 No errors are defined.

22312 **EXAMPLES**

22313 None.

22314 **APPLICATION USAGE**

22315 To ensure applications portability, especially across natural languages, only this function and  
22316 those listed in the SEE ALSO section should be used for classification of wide-character codes.

22317 **RATIONALE**

22318 None.

22319 **FUTURE DIRECTIONS**

22320 None.

22321 **SEE ALSO**

22322 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
22323 *iswpunct()*, *iswspace()*, *iswxdigit()*, *setlocale()*, the Base Definitions volume of  
22324 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>, the Base Definitions volume of  
22325 IEEE Std. 1003.1-200x, Chapter 7, Locale

22326 **CHANGE HISTORY**

22327 First released in Issue 4.

22328 **Issue 5**

22329 The following change has been made in this issue for alignment with  
22330 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22331 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
22332 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

22333 **Issue 6**

22334

The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22335 **NAME**

22336 iswxdigit — test for a hexadecimal digit wide-character code

22337 **SYNOPSIS**

22338 #include <wctype.h>

22339 int iswxdigit(wint\_t wc);

22340 **DESCRIPTION**

22341 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
22342 conflict between the requirements described here and the ISO C standard is unintentional. This  
22343 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22344 The *iswxdigit()* function shall test whether *wc* is a wide-character code representing a character  
22345 of class **xdigit** in the program's current locale; see the Base Definitions volume of  
22346 IEEE Std. 1003.1-200x, Chapter 7, Locale.

22347 In all cases *wc* is a **wint\_t**, the value of which the application shall ensure is a wide-character  
22348 code corresponding to a valid character in the current locale, or equal to the value of the macro  
22349 WEOF. If the argument has any other value, the behavior is undefined.

22350 **RETURN VALUE**

22351 The *iswxdigit()* function shall return non-zero if *wc* is a hexadecimal digit wide-character code;  
22352 otherwise, it shall return 0.

22353 **ERRORS**

22354 No errors are defined.

22355 **EXAMPLES**

22356 None.

22357 **APPLICATION USAGE**

22358 To ensure applications portability, especially across natural languages, only this function and  
22359 those listed in the SEE ALSO section should be used for classification of wide-character codes.

22360 **RATIONALE**

22361 None.

22362 **FUTURE DIRECTIONS**

22363 None.

22364 **SEE ALSO**

22365 *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswdigit()*, *iswgraph()*, *iswlower()*, *iswprint()*,  
22366 *iswpunct()*, *iswspace()*, *iswupper()*, *setlocale()*, the Base Definitions volume of  
22367 IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>

22368 **CHANGE HISTORY**

22369 First released in Issue 4.

22370 **Issue 5**

22371 The following change has been made in this issue for alignment with  
22372 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 22373 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
22374 now made visible by inclusion of the <wctype.h> header rather than <wchar.h>.

22375 **Issue 6**

22376 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22377 **NAME**

22378 isxdigit — test for a hexadecimal digit

22379 **SYNOPSIS**

22380 #include <ctype.h>

22381 int isxdigit(int c);

22382 **DESCRIPTION**

22383 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
22384 conflict between the requirements described here and the ISO C standard is unintentional. This  
22385 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22386 The *isxdigit()* function shall test whether *c* is a character of class **xdigit** in the program's current  
22387 locale; see the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale.

22388 In all cases *c* is an **int**, the value of which the application shall ensure is a character representable  
22389 as an **unsigned char** or equal to the value of the macro EOF. If the argument has any other value,  
22390 the behavior is undefined.

22391 **RETURN VALUE**

22392 The *isxdigit()* function shall return non-zero if *c* is a hexadecimal digit; otherwise, it shall return  
22393 0.

22394 **ERRORS**

22395 No errors are defined.

22396 **EXAMPLES**

22397 None.

22398 **APPLICATION USAGE**

22399 To ensure applications portability, especially across natural languages, only this function and  
22400 those listed in the SEE ALSO section should be used for character classification.

22401 **RATIONALE**

22402 None.

22403 **FUTURE DIRECTIONS**

22404 None.

22405 **SEE ALSO**

22406 *isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
22407 the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>

22408 **CHANGE HISTORY**

22409 First released in Issue 1. Derived from Issue 1 of the SVID.

22410 **Issue 4**

22411 The text of the DESCRIPTION is revised, although there are no functional differences between  
22412 this issue and Issue 3.

22413 **Issue 6**

22414 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22415 **NAME**

22416           j0, j1, jn — Bessel functions of the first kind

22417 **SYNOPSIS**

```
22418 XSI #include <math.h>
22419 double j0(double x);
22420 double j1(double x);
22421 double jn(int n, double x);
22422
```

22423 **DESCRIPTION**

22424           The *j0()*, *j1()*, and *jn()* functions shall compute Bessel functions of *x* of the first kind of orders 0, 1, and *n* respectively.

22426           An application wishing to check for error situations should set *errno* to 0 before calling *j0()*, *j1()*, or *jn()*. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

22428 **RETURN VALUE**

22429           Upon successful completion, *j0()*, *j1()*, and *jn()* shall return the relevant Bessel value of *x* of the first kind.

22431           If the *x* argument is too large in magnitude, 0 shall be returned and *errno* may be set to [ERANGE].

22433           If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

22434           If the correct result would cause underflow, 0 shall be returned and *errno* may be set to [ERANGE].

22436 **ERRORS**

22437           The *j0()*, *j1()*, and *jn()* functions may fail if:

22438           [EDOM]           The value of *x* is NaN.

22439           [ERANGE]        The value of *x* was too large in magnitude, or underflow occurred.

22440           No other errors shall occur.

22441 **EXAMPLES**

22442           None.

22443 **APPLICATION USAGE**

22444           None.

22445 **RATIONALE**

22446           None.

22447 **FUTURE DIRECTIONS**

22448           None.

22449 **SEE ALSO**

22450           *isnan()*, *y0()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

22451 **CHANGE HISTORY**

22452           First released in Issue 1. Derived from Issue 1 of the SVID.

22453 **Issue 4**

22454           References to *matherr()* are removed.

22455           The RETURN VALUE and ERRORS sections are substantially rewritten to rationalize error handling in the mathematics functions.

22457 **Issue 5**

22458

22459

The DESCRIPTION is updated to indicate how an application should check for an error. This text was previously published in the APPLICATION USAGE section.

22460 **NAME**

22461       jrand48 — generate a uniformly distributed pseudo-random long signed integer

22462 **SYNOPSIS**

22463 XSI       #include <stdlib.h>

22464       long jrand48(unsigned short xsubi[3]);

22465

22466 **DESCRIPTION**

22467       Refer to *drand48()*.

22468 **NAME**

22469 kill — send a signal to a process or a group of processes

22470 **SYNOPSIS**

22471 #include &lt;signal.h&gt;

22472 int kill(pid\_t pid, int sig);

22473 **DESCRIPTION**

22474 The *kill()* function shall send a signal to a process or a group of processes specified by *pid*. The  
 22475 signal to be sent is specified by *sig* and is either one from the list given in <**signal.h**> or 0. If *sig* is  
 22476 0 (the null signal), error checking is performed but no signal is actually sent. The null signal can  
 22477 be used to check the validity of *pid*.

22478 For a process to have permission to send a signal to a process designated by *pid*, unless the  
 22479 sending process has appropriate privileges, the application shall ensure that the real or effective  
 22480 user ID of the sending process matches the real or saved set-user-ID of the receiving process.

22481 If *pid* is greater than 0, *sig* shall be sent to the process whose process ID is equal to *pid*.

22482 If *pid* is 0, *sig* shall be sent to all processes (excluding an unspecified set of system processes)  
 22483 whose process group ID is equal to the process group ID of the sender, and for which the  
 22484 process has permission to send a signal.

22485 If *pid* is -1, *sig* shall be sent to all processes (excluding an unspecified set of system processes) for  
 22486 which the process has permission to send that signal.

22487 If *pid* is negative, but not -1, *sig* shall be sent to all processes (excluding an unspecified set of  
 22488 system processes) whose process group ID is equal to the absolute value of *pid*, and for which  
 22489 the process has permission to send a signal.

22490 If the value of *pid* causes *sig* to be generated for the sending process, and if *sig* is not blocked for  
 22491 the calling thread and if no other thread has *sig* unblocked or is waiting in a *sigwait()* function  
 22492 for *sig*, either *sig* or at least one pending unblocked signal shall be delivered to the sending  
 22493 thread before *kill()* returns.

22494 The user ID tests described above shall not be applied when sending SIGCONT to a process that  
 22495 is a member of the same session as the sending process.

22496 An implementation that provides extended security controls may impose further  
 22497 implementation-defined restrictions on the sending of signals, including the null signal. In  
 22498 particular, the system may deny the existence of some or all of the processes specified by *pid*.

22499 The *kill()* function is successful if the process has permission to send *sig* to any of the processes  
 22500 specified by *pid*. If *kill()* fails, no signal shall be sent.

22501 **RETURN VALUE**

22502 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
 22503 indicate the error.

22504 **ERRORS**

22505 The *kill()* function shall fail if:

22506 [EINVAL] The value of the *sig* argument is an invalid or unsupported signal number.

22507 [EPERM] The process does not have permission to send the signal to any receiving  
 22508 process.

22509 [ESRCH] No process or process group can be found corresponding to that specified by  
 22510 *pid*.

22511 **EXAMPLES**

22512 None.

22513 **APPLICATION USAGE**

22514 None.

22515 **RATIONALE**

22516 The semantics for permission checking for *kill()* differed between System V and most other  
22517 implementations, such as Version 7 or 4.3 BSD. The semantics chosen for this volume of  
22518 IEEE Std. 1003.1-200x agree with System V. Specifically, a set-user-ID process cannot protect  
22519 itself against signals (or at least not against SIGKILL) unless it changes its real user ID. This  
22520 choice allows the user who starts an application to send it signals even if it changes its effective  
22521 user ID. The other semantics give more power to an application that wants to protect itself from  
22522 the user who ran it.

22523 Some implementations provide semantic extensions to the *kill()* function when the absolute  
22524 value of *pid* is greater than some maximum, or otherwise special, value. Negative values are a  
22525 flag to *kill()*. Since most implementations return [ESRCH] in this case, this behavior is not  
22526 included in this volume of IEEE Std. 1003.1-200x, although a conforming implementation could  
22527 provide such an extension.

22528 The implementation-defined processes to which a signal cannot be sent may include the  
22529 scheduler or *init*.

22530 There was initially strong sentiment to specify that, if *pid* specifies that a signal be sent to the  
22531 calling process and that signal is not blocked, that signal would be delivered before *kill()*  
22532 returns. This would permit a process to call *kill()* and be guaranteed that the call never return.  
22533 However, historical implementations that provide only the *signal()* function make only the  
22534 weaker guarantee in this volume of IEEE Std. 1003.1-200x, because they only deliver one signal  
22535 each time a process enters the kernel. Modifications to such implementations to support the  
22536 *sigaction()* function generally require entry to the kernel following return from a signal-catching  
22537 function, in order to restore the signal mask. Such modifications have the effect of satisfying the  
22538 stronger requirement, at least when *sigaction()* is used, but not necessarily when *signal()* is used.  
22539 The developers of this volume of IEEE Std. 1003.1-200x considered making the stronger  
22540 requirement except when *signal()* is used, but felt this would be unnecessarily complex.  
22541 Implementors are encouraged to meet the stronger requirement whenever possible. In practice,  
22542 the weaker requirement is the same, except in the rare case when two signals arrive during a  
22543 very short window. This reasoning also applies to a similar requirement for *sigprocmask()*.

22544 In 4.2 BSD, the SIGCONT signal can be sent to any descendant process regardless of user-ID  
22545 security checks. This allows a job control shell to continue a job even if processes in the job have  
22546 altered their user IDs (as in the *su* command). In keeping with the addition of the concept of  
22547 sessions, similar functionality is provided by allowing the SIGCONT signal to be sent to any  
22548 process in the same session regardless of user ID security checks. This is less restrictive than BSD  
22549 in the sense that ancestor processes (in the same session) can now be the recipient. It is more  
22550 restrictive than BSD in the sense that descendant processes that form new sessions are now  
22551 subject to the user ID checks. A similar relaxation of security is not necessary for the other job  
22552 control signals since those signals are typically sent by the terminal driver in recognition of  
22553 special characters being typed; the terminal driver bypasses all security checks.

22554 In secure implementations, a process may be restricted from sending a signal to a process having  
22555 a different security label. In order to prevent the existence or nonexistence of a process from  
22556 being used as a covert channel, such processes should appear nonexistent to the sender; that is,  
22557 [ESRCH] should be returned, rather than [EPERM], if *pid* refers only to such processes.

22558 Existing implementations vary on the result of a *kill()* with *pid* indicating an inactive process (a  
 22559 terminated process that has not been waited for by its parent). Some indicate success on such a  
 22560 call (subject to permission checking), while others give an error of [ESRCH]. Since the definition  
 22561 of process lifetime in this volume of IEEE Std. 1003.1-200x covers inactive processes, the  
 22562 [ESRCH] error as described is inappropriate in this case. In particular, this means that an  
 22563 application cannot have a parent process check for termination of a particular child with *kill()*.  
 22564 (Usually this is done with the null signal; this can be done reliably with *waitpid()*.)

22565 There is some belief that the name *kill()* is misleading, since the function is not always intended  
 22566 to cause process termination. However, the name is common to all historical implementations,  
 22567 and any change would be in conflict with the goal of minimal changes to existing application  
 22568 code.

#### 22569 FUTURE DIRECTIONS

22570 None.

#### 22571 SEE ALSO

22572 *getpid()*, *raise()*, *setsid()*, *sigaction()*, *sigqueue()*, the Base Definitions volume of  
 22573 IEEE Std. 1003.1-200x, `<signal.h>`, `<sys/types.h>`

#### 22574 CHANGE HISTORY

22575 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 22576 Issue 4

22577 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on  
 22578 XSI-conformant systems.

22579 The DESCRIPTION is clarified in various places.

22580 The following change is incorporated for alignment with the FIPS requirements:

- 22581 • In the DESCRIPTION, the second paragraph is reworded to indicate that the saved set-user-  
 22582 ID of the calling process is checked in place of its effective user ID. This functionality is  
 22583 marked as an extension.

#### 22584 Issue 5

22585 The DESCRIPTION is updated for alignment with POSIX Threads Extension.

#### 22586 Issue 6

22587 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

22588 The following new requirements on POSIX implementations derive from alignment with the  
 22589 Single UNIX Specification:

- 22590 • In the DESCRIPTION, the second paragraph is reworded to indicate that the saved set-user-  
 22591 ID of the calling process is checked in place of its effective user ID. This is a FIPS  
 22592 requirement.
- 22593 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
 22594 required for conforming implementations of previous POSIX specifications, it was not  
 22595 required for UNIX applications.
- 22596 • The behavior when *pid* is `-1` is now specified. It was previously explicitly unspecified in the  
 22597 POSIX.1-1988 standard.

22598 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

22599 **NAME**

22600 killpg — send a signal to a process group

22601 **SYNOPSIS**

22602 XSI #include <signal.h>

22603 int killpg(pid\_t pgrp, int sig);

22604

22605 **DESCRIPTION**

22606 The *killpg()* function shall send the signal specified by *sig* to the process group specified by *pgrp*.

22607 If *pgrp* is greater than 1, *killpg(pgrp, sig)* shall be equivalent to *kill(-pgrp, sig)*. If *pgrp* is less than or  
22608 equal to 1, the behavior of *killpg()* is undefined.

22609 **RETURN VALUE**

22610 Refer to *kill()*.

22611 **ERRORS**

22612 Refer to *kill()*.

22613 **EXAMPLES**

22614 None.

22615 **APPLICATION USAGE**

22616 None.

22617 **RATIONALE**

22618 None.

22619 **FUTURE DIRECTIONS**

22620 None.

22621 **SEE ALSO**

22622 *getpgid()*, *getpid()*, *kill()*, *raise()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
22623 <signal.h>

22624 **CHANGE HISTORY**

22625 First released in Issue 4, Version 2.

22626 **Issue 5**

22627 Moved from X/OPEN UNIX extension to BASE.

22628 **NAME**

22629       l64a — convert a 32-bit integer to a radix-64 ASCII string

22630 **SYNOPSIS**

22631 xSI       #include &lt;stdlib.h&gt;

22632       char \*l64a(long value);

22633

22634 **DESCRIPTION**22635       Refer to *a64l()*.

22636 **NAME**

22637 labs, llabs — return a long integer absolute value

22638 **SYNOPSIS**

22639 #include <stdlib.h>

22640 long labs(long *i*);

22641 long long llabs(long long *i*);

22642 **DESCRIPTION**

22643 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
22644 conflict between the requirements described here and the ISO C standard is unintentional. This  
22645 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22646 These functions shall compute the absolute value of the **long** integer operand, *i*. If the result  
22647 cannot be represented, the behavior is undefined.

22648 **RETURN VALUE**

22649 These functions shall return the absolute value of the **long** integer operand.

22650 **ERRORS**

22651 No errors are defined.

22652 **EXAMPLES**

22653 None.

22654 **APPLICATION USAGE**

22655 None.

22656 **RATIONALE**

22657 None.

22658 **FUTURE DIRECTIONS**

22659 None.

22660 **SEE ALSO**

22661 *abs()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

22662 **CHANGE HISTORY**

22663 First released in Issue 4. Derived from the ISO C standard.

22664 **Issue 6**

22665 The *llabs()* function is added for alignment with the ISO/IEC 9899:1999 standard.

22666 **NAME**

22667 lchown — change the owner and group of a symbolic link

22668 **SYNOPSIS**22669 XSI `#include <unistd.h>`22670 `int lchown(const char *path, uid_t owner, gid_t group);`

22671

22672 **DESCRIPTION**

22673 The *lchown()* function shall have the same effect as *chown()* except in the case where the named  
 22674 file is a symbolic link. In this case, *lchown()* shall change the ownership of the symbolic link file  
 22675 itself, while *chown()* changes the ownership of the file or directory to which the symbolic link  
 22676 refers.

22677 **RETURN VALUE**

22678 Upon successful completion, *lchown()* shall return 0. Otherwise, it shall return  $-1$  and set *errno* to  
 22679 indicate an error.

22680 **ERRORS**22681 The *lchown()* function shall fail if:

- |       |                |                                                                                              |
|-------|----------------|----------------------------------------------------------------------------------------------|
| 22682 | [EACCES]       | Search permission is denied on a component of the path prefix of <i>path</i> .               |
| 22683 | [EINVAL]       | The owner or group ID is not a value supported by the implementation.                        |
| 22684 | [ELOOP]        | A loop exists in symbolic links encountered during resolution of the <i>path</i>             |
| 22685 |                | argument.                                                                                    |
| 22686 | [ENAMETOOLONG] |                                                                                              |
| 22687 |                | The length of a path name exceeds {PATH_MAX} or a path name component                        |
| 22688 |                | is longer than {NAME_MAX}.                                                                   |
| 22689 | [ENOENT]       | A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string. |
| 22690 | [ENOTDIR]      | A component of the path prefix of <i>path</i> is not a directory.                            |
| 22691 | [EOPNOTSUPP]   | The <i>path</i> argument names a symbolic link and the implementation does not               |
| 22692 |                | support setting the owner or group of a symbolic link.                                       |
| 22693 | [EPERM]        | The effective user ID does not match the owner of the file and the process                   |
| 22694 |                | does not have appropriate privileges.                                                        |
| 22695 | [EROFS]        | The file resides on a read-only file system.                                                 |

22696 The *lchown()* function may fail if:

- |       |                |                                                                         |
|-------|----------------|-------------------------------------------------------------------------|
| 22697 | [EIO]          | An I/O error occurred while reading or writing to the file system.      |
| 22698 | [EINTR]        | A signal was caught during execution of the function.                   |
| 22699 | [ELOOP]        | More than {SYMLOOP_MAX} symbolic links were encountered during          |
| 22700 |                | resolution of the <i>path</i> argument.                                 |
| 22701 | [ENAMETOOLONG] |                                                                         |
| 22702 |                | Path name resolution of a symbolic link produced an intermediate result |
| 22703 |                | whose length exceeds {PATH_MAX}.                                        |

22704 **EXAMPLES**22705 **Changing the Current Owner of a File**

22706 The following example shows how to change the ownership of the symbolic link named  
22707 **/modules/pass1** to the user ID associated with “jones” and the group ID associated with “cnd”.

22708 The numeric value for the user ID is obtained by using the *getpwnam()* function. The numeric  
22709 value for the group ID is obtained by using the *getgrnam()* function.

```
22710 #include <sys/types.h>
22711 #include <unistd.h>
22712 #include <pwd.h>
22713 #include <grp.h>

22714 struct passwd *pwd;
22715 struct group *grp;
22716 char *path = "/modules/pass1";
22717 ...
22718 pwd = getpwnam("jones");
22719 grp = getgrnam("cnd");
22720 lchown(path, pwd->pw_uid, grp->gr_gid);
```

22721 **APPLICATION USAGE**

22722 None.

22723 **RATIONALE**

22724 None.

22725 **FUTURE DIRECTIONS**

22726 None.

22727 **SEE ALSO**

22728 *chown()*, *symlink()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<unistd.h>**

22729 **CHANGE HISTORY**

22730 First released in Issue 4, Version 2.

22731 **Issue 5**

22732 Moved from X/OPEN UNIX extension to BASE.

22733 **Issue 6**

22734 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
22735 [ELOOP] error condition is added.

22736 **NAME**

22737        lcong48 — seed a uniformly distributed pseudo-random signed long integer generator

22738 **SYNOPSIS**

22739 xSI       #include <stdlib.h>

22740       void lcong48(unsigned short param[7]);

22741

22742 **DESCRIPTION**

22743       Refer to *drand48()*.

22744 **NAME**

22745 ldexp, ldexpf, ldexpl — load exponent of a floating point number

22746 **SYNOPSIS**

22747 #include <math.h>

22748 double ldexp(double x, int exp);

22749 float ldexpf(float x, int exp);

22750 long double ldexpl(long double x, int exp);

22751 **DESCRIPTION**

22752 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 22753 conflict between the requirements described here and the ISO C standard is unintentional. This  
 22754 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22755 These functions shall compute the quantity  $x * 2^{exp}$ .

22756 An application wishing to check for error situations should set *errno* to 0 before calling *ldexp()*.

22757 If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

22758 **RETURN VALUE**

22759 Upon successful completion, these functions shall return *x* multiplied by 2, raised to the power  
 22760 *exp*.

22761 **XSI** If the value of *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

22762 If *ldexp()* would cause overflow,  $\pm$ HUGE\_VAL shall be returned (according to the sign of *x*), and  
 22763 *errno* shall be set to [ERANGE].

22764 If *ldexp()* would cause underflow, 0 shall be returned and *errno* may be set to [ERANGE].

22765 **ERRORS**

22766 These functions shall fail if:

22767 [ERANGE] The value to be returned would have caused overflow.

22768 These functions may fail if:

22769 **XSI** [EDOM] The argument *x* is NaN.

22770 [ERANGE] The value to be returned would have caused underflow.

22771 No other errors shall occur.

22772 **EXAMPLES**

22773 None.

22774 **APPLICATION USAGE**

22775 None.

22776 **RATIONALE**

22777 None.

22778 **FUTURE DIRECTIONS**

22779 None.

22780 **SEE ALSO**

22781 *frexp()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

22782 **CHANGE HISTORY**

22783 First released in Issue 1. Derived from Issue 1 of the SVID.

22784 **Issue 4**

22785 References to *matherr()* are removed.

22786 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
22787 ISO C standard and to rationalize error handling in the mathematics functions.

22788 The return value specified for [EDOM] is marked as an extension.

22789 **Issue 5**

22790 The DESCRIPTION is updated to indicate how an application should check for an error. This  
22791 text was previously published in the APPLICATION USAGE section.

22792 **Issue 6**

22793 The *ldexpf()* and *ldexpl()* functions are added for alignment with the ISO/IEC 9899:1999  
22794 standard.

22795 **NAME**

22796 ldiv, lldiv — compute quotient and remainder of a long division

22797 **SYNOPSIS**

22798 #include <stdlib.h>

22799 ldiv\_t ldiv(long *numer*, long *denom*);

22800 lldiv\_t lldiv(long long *numer*, long long *denom*);

22801 **DESCRIPTION**

22802 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
22803 conflict between the requirements described here and the ISO C standard is unintentional. This  
22804 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

22805 These functions shall compute the quotient and remainder of the division of the numerator  
22806 *numer* by the denominator *denom*. If the division is inexact, the resulting quotient is the **long**  
22807 integer of lesser magnitude that is the nearest to the algebraic quotient. If the result cannot be  
22808 represented, the behavior is undefined; otherwise, *quot* \* *denom* + *rem* shall equal *numer*.

22809 **RETURN VALUE**

22810 These functions shall return a structure of type **ldiv\_t**, comprising both the quotient and the  
22811 remainder. The structure includes the following members, in any order:

22812 long quot; /\* Quotient \*/

22813 long rem; /\* Remainder \*/

22814 **ERRORS**

22815 No errors are defined.

22816 **EXAMPLES**

22817 None.

22818 **APPLICATION USAGE**

22819 None.

22820 **RATIONALE**

22821 None.

22822 **FUTURE DIRECTIONS**

22823 None.

22824 **SEE ALSO**

22825 *div()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

22826 **CHANGE HISTORY**

22827 First released in Issue 4. Derived from the ISO C standard.

22828 **Issue 6**

22829 The *lldiv()* function is added for alignment with the ISO/IEC 9899:1999 standard.

22830 **NAME**

22831 lfind — find entry in a linear search table

22832 **SYNOPSIS**

22833 XSI #include &lt;search.h&gt;

22834 void \*lfind(const void \*key, const void \*base, size\_t \*nelp,  
22835 size\_t width, int (\*compar)(const void \*, const void \*));

22836

22837 **DESCRIPTION**22838 Refer to *lsearch()*.

22839 **NAME**

22840 lgamma, lgammaf, lgammal — log gamma function

22841 **SYNOPSIS**

```
22842 xSI #include <math.h>
22843 double lgamma(double x);
22844 float lgammaf(float x);
22845 long double lgammal(long double x);
22846 extern int signgam;
22847
```

22848 **DESCRIPTION**

22849 These functions shall compute  $\log_e|\Gamma(x)|$  where  $\Gamma(x)$  is defined as  $\int_0^{\infty} e^{-t} t^{x-1} dt$ . The sign of  
 22850  $\Gamma(x)$  is returned in the external integer *signgam*. The argument *x* need not be a non-positive  
 22851 integer ( $\Gamma(x)$  is defined over the reals, except the non-positive integers).  
 22852

22853 An application wishing to check for error situations should set *errno* to 0 before calling *lgamma()*.  
 22854 If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

22855 These functions need not be reentrant. A function that is not required to be reentrant is not  
 22856 required to be thread-safe.

22857 **RETURN VALUE**

22858 Upon successful completion, these functions shall return the logarithmic gamma of *x*.

22859 If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

22860 If *x* is a non-positive integer, either HUGE\_VAL or NaN shall be returned and *errno* may be set to  
 22861 [ERANGE].

22862 If the correct value would cause overflow, *lgamma()* shall return HUGE\_VAL and may set *errno*  
 22863 to [ERANGE].

22864 If the correct value would cause underflow, *lgamma()* shall return 0 and may set *errno* to  
 22865 [ERANGE].

22866 **ERRORS**

22867 These functions may fail if:

22868 [EDOM] The value of *x* is NaN.

22869 [ERANGE] The value of *x* is a non-positive integer, or the value to be returned would  
 22870 have caused overflow or underflow.

22871 No other errors shall occur.

22872 **EXAMPLES**

22873 None.

22874 **APPLICATION USAGE**

22875 None.

22876 **RATIONALE**

22877 None.

22878 **FUTURE DIRECTIONS**

22879 None.

22880 **SEE ALSO**

22881 *exp()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

22882 **CHANGE HISTORY**

22883 First released in Issue 3.

22884 **Issue 4**

22885 This page no longer points to *gamma()*, but contains all information relating to *lgamma()*.

22886 The RETURN VALUE and ERRORS sections are substantially rewritten to rationalize error handling in the mathematics functions.

22888 **Issue 5**

22889 The DESCRIPTION is updated to indicate how an application should check for an error. This text was previously published in the APPLICATION USAGE section.

22891 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

22892 **Issue 6**

22893 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 22894 • The *lgammaf()* and *lgammal()* functions are added.
  - 22895 • The RETURN VALUE and ERRORS sections are updated so that when *x* is a non-positive integer, *errno* may be set to [ERANGE]; previously, this was [EDOM].
- 22896

## 22897 NAME

22898 link — link to a file

## 22899 SYNOPSIS

22900 #include &lt;unistd.h&gt;

22901 int link(const char \*path1, const char \*path2);

## 22902 DESCRIPTION

22903 The *link()* function shall create a new link (directory entry) for the existing file, *path1*.22904 The *path1* argument points to a path name naming an existing file. The *path2* argument points to  
22905 a path name naming the new directory entry to be created. The *link()* function shall atomically  
22906 create a new link for the existing file and the link count of the file shall be incremented by one.22907 If *path1* names a directory, *link()* shall fail unless the process has appropriate privileges and the  
22908 implementation supports using *link()* on directories.22909 Upon successful completion, *link()* shall mark for update the *st\_ctime* field of the file. Also, the  
22910 *st\_ctime* and *st\_mtime* fields of the directory that contains the new entry shall be marked for  
22911 update.22912 If *link()* fails, no link shall be created and the link count of the file shall remain unchanged.22913 The implementation may require that the calling process has permission to access the existing  
22914 file.

## 22915 RETURN VALUE

22916 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
22917 indicate the error.

## 22918 ERRORS

22919 The *link()* function shall fail if:22920 [EACCES] A component of either path prefix denies search permission, or the requested  
22921 link requires writing in a directory that denies write permission, or the calling  
22922 process does not have permission to access the existing file and this is  
22923 required by the implementation.22924 [EEXIST] The *path2* argument resolves to an existing file or refers to a symbolic link.22925 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path1* or  
22926 *path2* argument.22927 [EMLINK] The number of links to the file named by *path1* would exceed {LINK\_MAX}.

22928 [ENAMETOOLONG]

22929 The length of the *path1* or *path2* argument exceeds {PATH\_MAX} or a path  
22930 name component is longer than {NAME\_MAX}.22931 [ENOENT] A component of either path prefix does not exist; the file named by *path1* does  
22932 not exist; or *path1* or *path2* points to an empty string.

22933 [ENOSPC] The directory to contain the link cannot be extended.

22934 [ENOTDIR] A component of either path prefix is not a directory.

22935 [EPERM] The file named by *path1* is a directory and either the calling process does not  
22936 have appropriate privileges or the implementation prohibits using *link()* on  
22937 directories.

|       |                |                                                                                                                                                                             |
|-------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 22938 | [EROFS]        | The requested link requires writing in a directory on a read-only file system.                                                                                              |
| 22939 | [EXDEV]        | The link named by <i>path2</i> and the file named by <i>path1</i> are on different file systems and the implementation does not support links between file systems,         |
| 22940 | XSR            | or <i>path1</i> refers to a named STREAM.                                                                                                                                   |
| 22941 |                |                                                                                                                                                                             |
| 22942 |                | The <i>link()</i> function may fail if:                                                                                                                                     |
| 22943 | [ELOOP]        | More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path1</i> or <i>path2</i> argument.                                                     |
| 22944 |                |                                                                                                                                                                             |
| 22945 | [ENAMETOOLONG] |                                                                                                                                                                             |
| 22946 |                | As a result of encountering a symbolic link in resolution of the <i>path1</i> or <i>path2</i> argument, the length of the substituted path name string exceeded {PATH_MAX}. |
| 22947 |                |                                                                                                                                                                             |
| 22948 |                |                                                                                                                                                                             |

## 22949 EXAMPLES

### 22950 Creating a Link to a File

22951 The following example shows how to create a link to a file named `/home/cnd/mod1` by creating a  
22952 new directory entry named `/modules/pass1`.

```
22953 #include <unistd.h>
22954
22954 char *path1 = "/home/cnd/mod1";
22955 char *path2 = "/modules/pass1";
22956 int status;
22957 ...
22958 status = link (path1, path2);
```

### 22959 Creating a Link to a File Within a Program

22960 In the following program example, the *link()* function is used to link the `/etc/passwd` file  
22961 (defined as `PASSWDFILE`) to a file named `/etc/opasswd` (defined as `SAVEFILE`), which is used  
22962 to save the current password file. Then, after removing the current password file (defined as  
22963 `PASSWDFILE`), the new password file is saved as the current password file using the *link()*  
22964 function again.

```
22965 #include <unistd.h>
22966
22966 #define LOCKFILE "/etc/ptmp"
22967 #define PASSWDFILE "/etc/passwd"
22968 #define SAVEFILE "/etc/opasswd"
22969 ...
22970 /* Save current password file */
22971 link (PASSWDFILE, SAVEFILE);
22972
22972 /* Remove current password file. */
22973 unlink (PASSWDFILE);
22974
22974 /* Save new password file as current password file. */
22975 link (LOCKFILE, PASSWDFILE);
```

### 22976 APPLICATION USAGE

22977 Some implementations do allow links between file systems.

22978 **RATIONALE**

22979 Linking to a directory is restricted to the superuser in most historical implementations because  
22980 this capability may produce loops in the file hierarchy or otherwise corrupt the file system. This  
22981 volume of IEEE Std. 1003.1-200x continues that philosophy by prohibiting *link()* and *unlink()*  
22982 from doing this. Other functions could do it if the implementor designed such an extension.

22983 Some historical implementations allow linking of files on different file systems. Wording was  
22984 added to explicitly allow this optional behavior.

22985 The exception for cross-file system links is intended to apply only to links that are  
22986 programmatically indistinguishable from “hard” links.

22987 **FUTURE DIRECTIONS**

22988 None.

22989 **SEE ALSO**

22990 *symlink()*, *unlink()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

22991 **CHANGE HISTORY**

22992 First released in Issue 1. Derived from Issue 1 of the SVID.

22993 **Issue 4**

22994 The <**unistd.h**> header is added to the SYNOPSIS section.

22995 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 22996 • The type of arguments *path1* and *path2* are changed from **char\*** to **const char\***.

22997 The following change is incorporated for alignment with the FIPS requirements:

- 22998 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
22999 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
23000 an extension.

23001 **Issue 4, Version 2**

23002 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 23003 • The [ELOOP] error is returned if too many symbolic links are encountered during path name  
23004 resolution.
- 23005 • The [EXDEV] error may also be returned if *path1* refers to a named STREAM.
- 23006 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an  
23007 intermediate result of path name resolution of a symbolic link.

23008 **Issue 6**

23009 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 23010 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.  
23011 This is since behavior may vary from one file system to another.

23012 The following new requirements on POSIX implementations derive from alignment with the  
23013 Single UNIX Specification:

- 23014 • The [ELOOP] mandatory error condition is added.
- 23015 • A second [ENAMETOOLONG] is added as an optional error condition.

23016 The following changes were made to align with the IEEE P1003.1a draft standard:

- 23017 • An explanation is added of action when *path2* refers to a symbolic link.

23018

- The [ELOOP] optional error condition is added.

## 23019 NAME

23020 lio\_listio — list directed I/O (**REALTIME**)

## 23021 SYNOPSIS

23022 AIO #include &lt;aio.h&gt;

```
23023 int lio_listio(int mode, struct aiocb *restrict const list[restrict],
23024 int nent, struct sigevent *restrict sig);
23025
```

## 23026 DESCRIPTION

23027 The *lio\_listio()* function allows the calling process to initiate a list of I/O requests with a single  
23028 function call.

23029 The *mode* argument takes one of the values LIO\_WAIT or LIO\_NOWAIT declared in <aio.h> and  
23030 determines whether the function returns when the I/O operations have been completed, or as  
23031 soon as the operations have been queued. If the *mode* argument is LIO\_WAIT, the function waits  
23032 until all I/O is complete and the *sig* argument is ignored.

23033 If the *mode* argument is LIO\_NOWAIT, the function shall return immediately, and asynchronous  
23034 notification shall occur, according to the *sig* argument, when all the I/O operations complete. If  
23035 *sig* is NULL, then no asynchronous notification shall occur. If *sig* is not NULL, asynchronous  
23036 notification occurs as specified in Section 2.4.1 (on page 528) when all the requests in *list* have  
23037 completed.

23038 The I/O requests enumerated by *list* are submitted in an unspecified order.

23039 The *list* argument is an array of pointers to **aiocb** structures. The array contains *nent* elements.  
23040 The array may contain NULL elements, which shall be ignored.

23041 The *aio\_lio\_opcode* field of each **aiocb** structure specifies the operation to be performed. The  
23042 supported operations are LIO\_READ, LIO\_WRITE, and LIO\_NOP; these symbols are defined in  
23043 <aio.h>. The LIO\_NOP operation causes the list entry to be ignored. If the *aio\_lio\_opcode*  
23044 element is equal to LIO\_READ, then an I/O operation is submitted as if by a call to *aio\_read()*  
23045 with the *aiocbp* equal to the address of the **aiocb** structure. If the *aio\_lio\_opcode* element is equal  
23046 to LIO\_WRITE, then an I/O operation is submitted as if by a call to *aio\_write()* with the *aiocbp*  
23047 equal to the address of the **aiocb** structure.

23048 The *aio\_fildes* member specifies the file descriptor on which the operation is to be performed.

23049 The *aio\_buf* member specifies the address of the buffer to or from which the data is transferred.

23050 The *aio\_nbytes* member specifies the number of bytes of data to be transferred.

23051 The members of the **aiocb** structure further describe the I/O operation to be performed, in a  
23052 manner identical to that of the corresponding **aiocb** structure when used by the *aio\_read()* and  
23053 *aio\_write()* functions.

23054 The *nent* argument specifies how many elements are members of the list; that is, the length of the  
23055 array.

23056 The behavior of this function is altered according to the definitions of synchronized I/O data  
23057 integrity completion and synchronized I/O file integrity completion if synchronized I/O is  
23058 enabled on the file associated with *aio\_fildes*.

23059 For regular files, no data transfer shall occur past the offset maximum established in the open  
23060 file description associated with *aiocbp->aio\_fildes*.

## 23061 RETURN VALUE

23062 If the *mode* argument has the value LIO\_NOWAIT, the *lio\_listio()* function shall return the value  
 23063 zero if the I/O operations are successfully queued; otherwise, the function shall return the value  
 23064  $-1$  and set *errno* to indicate the error.

23065 If the *mode* argument has the value LIO\_WAIT, the *lio\_listio()* function shall return the value  
 23066 zero when all the indicated I/O has completed successfully. Otherwise, *lio\_listio()* shall return a  
 23067 value of  $-1$  and set *errno* to indicate the error.

23068 In either case, the return value only indicates the success or failure of the *lio\_listio()* call itself,  
 23069 not the status of the individual I/O requests. In some cases one or more of the I/O requests  
 23070 contained in the list may fail. Failure of an individual request does not prevent completion of  
 23071 any other individual request. To determine the outcome of each I/O request, the application  
 23072 shall examine the error status associated with each **aiocb** control block. The error statuses so  
 23073 returned are identical to those returned as the result of an *aio\_read()* or *aio\_write()* function.

## 23074 ERRORS

23075 The *lio\_listio()* function shall fail if:

23076 [EAGAIN] The resources necessary to queue all the I/O requests were not available. The  
 23077 application may check the error status for each **aiocb** to determine the  
 23078 individual request(s) that failed.

23079 [EAGAIN] The number of entries indicated by *nent* would cause the system-wide limit  
 23080 {AIO\_MAX} to be exceeded.

23081 [EINVAL] The *mode* argument is not a proper value, or the value of *nent* was greater than  
 23082 {AIO\_LISTIO\_MAX}.

23083 [EINTR] A signal was delivered while waiting for all I/O requests to complete during  
 23084 an LIO\_WAIT operation. Note that, since each I/O operation invoked by  
 23085 *lio\_listio()* may possibly provoke a signal when it completes, this error return  
 23086 may be caused by the completion of one (or more) of the very I/O operations  
 23087 being awaited. Outstanding I/O requests are not canceled, and the application  
 23088 shall examine each list element to determine whether the request was  
 23089 initiated, canceled, or completed.

23090 [EIO] One or more of the individual I/O operations failed. The application may  
 23091 check the error status for each **aiocb** structure to determine the individual  
 23092 request(s) that failed.

23093 In addition to the errors returned by the *lio\_listio()* function, if the *lio\_listio()* function succeeds  
 23094 or fails with errors of [EAGAIN], [EINTR], or [EIO], then some of the I/O specified by the list  
 23095 may have been initiated. If the *lio\_listio()* function fails with an error code other than [EAGAIN],  
 23096 [EINTR], or [EIO], no operations from the list shall have been initiated. The I/O operation  
 23097 indicated by each list element can encounter errors specific to the individual read or write  
 23098 function being performed. In this event, the error status for each **aiocb** control block contains the  
 23099 associated error code. The error codes that can be set are the same as would be set by a *read()* or  
 23100 *write()* function, with the following additional error codes possible:

23101 [EAGAIN] The requested I/O operation was not queued due to resource limitations.

23102 [ECANCELED] The requested I/O was canceled before the I/O completed due to an explicit  
 23103 *aio\_cancel()* request.

23104 [EFBIG] The *aiocbp->aio\_lio\_opcode* is LIO\_WRITE, the file is a regular file, *aiocbp->*  
 23105 *aio\_nbytes* is greater than 0, and the *aiocbp->aio\_offset* is greater than or equal  
 23106 to the offset maximum in the open file description associated with *aiocbp-*

23107 >*aio\_fildes*.

23108 [EINPROGRESS] The requested I/O is in progress.

23109 [EOVERFLOW] The *aiochp->aio\_lio\_opcode* is LIO\_READ, the file is a regular file, *aiochp->aio\_nbytes* is greater than 0, and the *aiochp->aio\_offset* is before the end-of-file and is greater than or equal to the offset maximum in the open file description associated with *aiochp->aio\_fildes*.

#### 23113 EXAMPLES

23114 None.

#### 23115 APPLICATION USAGE

23116 None.

#### 23117 RATIONALE

23118 Although it may appear that there are inconsistencies in the specified circumstances for error codes, the [EIO] error condition applies when any circumstance relating to an individual operation makes that operation fail. This might be due to a badly formulated request (for example, the *aio\_lio\_opcode* field is invalid, and *aio\_error()* returns [EINVAL]) or might arise from application behavior (for example, the file descriptor is closed before the operation is initiated, and *aio\_error()* returns [EBADF]).

23124 The limitation on the set of error codes returned when operations from the list shall have been initiated enables applications to know when operations have been started and whether *aio\_error()* is valid for a specific operation.

#### 23127 FUTURE DIRECTIONS

23128 None.

#### 23129 SEE ALSO

23130 *aio\_read()*, *aio\_write()*, *aio\_error()*, *aio\_return()*, *aio\_cancel()*, *close()*, *exec*, *exit()*, *fork()*, *lseek()*, *read()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**aio.h**>

#### 23132 CHANGE HISTORY

23133 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

23134 Large File Summit extensions are added.

#### 23135 Issue 6

23136 The [ENOSYS] error condition has been removed as stubs need not be provided if an implementation does not support the Asynchronous Input and Output option.

23138 The *lio\_listio()* function is marked as part of the Asynchronous Input and Output option.

23139 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

23141 • In the DESCRIPTION, text is added to indicate that for regular files no data transfer occurs past the offset maximum established in the open file description associated with *aiochp->aio\_fildes*. This change is to support large files.

23144 • The [EBIG] and [EOVERFLOW] error conditions are defined. This change is to support large files.

23146 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

23147 The **restrict** keyword is added to the *lio\_listio()* prototype for alignment with the ISO/IEC 9899:1999 standard.

23149 **NAME**

23150 `listen` — listen for socket connections and limit the queue of incoming connections

23151 **SYNOPSIS**

23152 `#include <sys/socket.h>`

23153 `int listen(int socket, int backlog);`

23154 **DESCRIPTION**

23155 The `listen()` function shall mark a connection-mode socket, specified by the `socket` argument, as  
 23156 accepting connections.

23157 The `backlog` argument provides a hint to the implementation which the implementation shall use  
 23158 to limit the number of outstanding connections in the socket's listen queue. Implementations  
 23159 may impose a limit on `backlog` and silently reduce the specified value. Normally, a larger `backlog`  
 23160 argument value shall result in a larger or equal length of the listen queue. Implementations shall  
 23161 support values of `backlog` up to `SOMAXCONN`, defined in `<sys/socket.h>`.

23162 The implementation may include incomplete connections in its listen queue. The limits on the  
 23163 number of incomplete connections and completed connections queued may be different.

23164 The implementation may have an upper limit on the length of the listen queue—either global or  
 23165 per accepting socket. If `backlog` exceeds this limit, the length of the listen queue is set to the limit.

23166 If `listen()` is called with a `backlog` argument value that is less than 0, the function behaves as if it  
 23167 had been called with a `backlog` argument value of 0.

23168 A `backlog` argument of 0 may allow the socket to accept connections, in which case the length of  
 23169 the listen queue may be set to an implementation-defined minimum value.

23170 The socket in use may require the process to have appropriate privileges to use the `listen()`  
 23171 function.

23172 **RETURN VALUE**

23173 Upon successful completions, `listen()` shall return 0; otherwise, `-1` shall be returned and `errno` set  
 23174 to indicate the error.

23175 **ERRORS**

23176 The `listen()` function shall fail if:

23177 [EBADF] The `socket` argument is not a valid file descriptor.

23178 [EDESTADDRREQ]

23179 The socket is not bound to a local address, and the protocol does not support  
 23180 listening on an unbound socket.

23181 [EINVAL] The `socket` is already connected.

23182 [ENOTSOCK] The `socket` argument does not refer to a socket.

23183 [EOPNOTSUPP] The socket protocol does not support `listen()`.

23184 The `listen()` function may fail if:

23185 [EACCES] The calling process does not have the appropriate privileges.

23186 [EINVAL] The `socket` has been shut down.

23187 [ENOBUFS] Insufficient resources are available in the system to complete the call.

23188 **EXAMPLES**

23189           None.

23190 **APPLICATION USAGE**

23191           None.

23192 **RATIONALE**

23193           None.

23194 **FUTURE DIRECTIONS**

23195           None.

23196 **SEE ALSO**

23197           *accept()*, *connect()*, *socket()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/socket.h> |

23198 **CHANGE HISTORY**

23199           First released in Issue 6. Derived from the XNS, Issue 5.2 specification. |

23200           The DESCRIPTION is updated to describe the relationship of SOMAXCONN and the *backlog* |

23201           argument.

23202 **NAME**

23203 llrint, llrintf, llrintl, lrint, lrintf, lrintl — round to nearest integer value using current rounding  
23204 direction

23205 **SYNOPSIS**

```
23206 #include <math.h>

23207 long long llrint(double x);
23208 long long llrintf(float x);
23209 long long llrintl(long double x);
23210 long lrint(double x);
23211 long lrintf(float x);
23212 long lrintl(long double x);
```

23213 **DESCRIPTION**

23214 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
23215 conflict between the requirements described here and the ISO C standard is unintentional. This  
23216 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

23217 These functions shall round their argument to the nearest integer value, rounding according to  
23218 the current rounding direction.

23219 An application wishing to check for error situations should set *errno* to 0 before calling these  
23220 functions. If *errno* is non-zero on return, an error has occurred.

23221 **RETURN VALUE**

23222 Upon successful completion, these functions shall return the rounded integer value.

23223 If the rounded value is outside the range of the return type, the numeric result is unspecified.

23224 If the magnitude of *x* is too large, the numeric result is unspecified and *errno* may be set to  
23225 [ERANGE].

23226 **ERRORS**

23227 These functions may fail if:

23228 [ERANGE] The magnitude of *x* is too large.

23229 **EXAMPLES**

23230 None.

23231 **APPLICATION USAGE**

23232 None.

23233 **RATIONALE**

23234 None.

23235 **FUTURE DIRECTIONS**

23236 None.

23237 **SEE ALSO**

23238 The Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

23239 **CHANGE HISTORY**

23240 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

23241 **NAME**

23242 llround, llroundf, llroundl, lround, lroundf, lroundl — round to nearest integer value

23243 **SYNOPSIS**

23244 #include <math.h>

23245 long long llround(double x);

23246 long long llroundf(float x);

23247 long long llroundl(long double x);

23248 long lround(double x);

23249 long lroundf(float x);

23250 long lroundl(long double x);

23251 **DESCRIPTION**

23252 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
23253 conflict between the requirements described here and the ISO C standard is unintentional. This  
23254 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

23255 These functions shall round their argument to the nearest integer value, rounding halfway cases  
23256 away from zero, regardless of the current rounding direction.

23257 An application wishing to check for error situations should set *errno* to 0 before calling these  
23258 functions. If *errno* is non-zero on return, an error has occurred.

23259 **RETURN VALUE**

23260 Upon successful completion, these functions shall return the rounded integer value.

23261 If the rounded value is outside the range of the return type, the numeric result is unspecified.

23262 If the magnitude of *x* is too large, the numeric result is unspecified and *errno* may be set to  
23263 [ERANGE].

23264 **ERRORS**

23265 These functions may fail if:

23266 [ERANGE] The magnitude of *x* is too large.

23267 **EXAMPLES**

23268 None.

23269 **APPLICATION USAGE**

23270 None.

23271 **RATIONALE**

23272 None.

23273 **FUTURE DIRECTIONS**

23274 None.

23275 **SEE ALSO**

23276 The Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

23277 **CHANGE HISTORY**

23278 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

23279 **NAME**

23280 localeconv — return locale-specific information

23281 **SYNOPSIS**

23282 #include &lt;locale.h&gt;

23283 struct lconv \*localeconv(void);

23284 **DESCRIPTION**

23285 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 23286 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23287 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

23288 The *localeconv()* function shall set the components of an object with the type **struct lconv** with  
 23289 the values appropriate for the formatting of numeric quantities (monetary and otherwise)  
 23290 according to the rules of the current locale.

23291 The members of the structure with type **char\*** are pointers to strings, any of which (except  
 23292 **decimal\_point**) can point to " ", to indicate that the value is not available in the current locale or  
 23293 is of zero length. The members with type **char** are non-negative numbers, any of which can be  
 23294 {CHAR\_MAX} to indicate that the value is not available in the current locale.

23295 The members include the following:

23296 **char \*decimal\_point**

23297 The radix character used to format non-monetary quantities.

23298 **char \*thousands\_sep**

23299 The character used to separate groups of digits before the decimal-point character in  
 23300 formatted non-monetary quantities.

23301 **char \*grouping**

23302 A string whose elements taken as one-byte integer values indicate the size of each group of  
 23303 digits in formatted non-monetary quantities.

23304 **char \*int\_curr\_symbol**

23305 The international currency symbol applicable to the current locale. The first three  
 23306 characters contain the alphabetic international currency symbol in accordance with those  
 23307 specified in the ISO 4217:1995 standard. The fourth character (immediately preceding the  
 23308 null byte) is the character used to separate the international currency symbol from the  
 23309 monetary quantity.

23310 **char \*currency\_symbol**

23311 The local currency symbol applicable to the current locale.

23312 **char \*mon\_decimal\_point**

23313 The radix character used to format monetary quantities.

23314 **char \*mon\_thousands\_sep**

23315 The separator for groups of digits before the decimal-point in formatted monetary  
 23316 quantities.

23317 **char \*mon\_grouping**

23318 A string whose elements taken as one-byte integer values indicate the size of each group of  
 23319 digits in formatted monetary quantities.

23320 **char \*positive\_sign**

23321 The string used to indicate a non-negative valued formatted monetary quantity.

|           |                                                                                                          |
|-----------|----------------------------------------------------------------------------------------------------------|
| 23322     | <b>char *negative_sign</b>                                                                               |
| 23323     | The string used to indicate a negative valued formatted monetary quantity.                               |
| 23324     | <b>char int_frac_digits</b>                                                                              |
| 23325     | The number of fractional digits (those after the decimal-point) to be displayed in an                    |
| 23326     | internationally formatted monetary quantity.                                                             |
| 23327     | <b>char frac_digits</b>                                                                                  |
| 23328     | The number of fractional digits (those after the decimal-point) to be displayed in a                     |
| 23329     | formatted monetary quantity.                                                                             |
| 23330     | <b>char p_cs_precedes</b>                                                                                |
| 23331     | Set to 1 if the <b>currency_symbol</b> or <b>int_curr_symbol</b> precedes the value for a non-negative   |
| 23332     | formatted monetary quantity. Set to 0 if the symbol succeeds the value.                                  |
| 23333     | <b>char p_sep_by_space</b>                                                                               |
| 23334     | Set to 0 if no space separates the <b>currency_symbol</b> or <b>int_curr_symbol</b> from the value for a |
| 23335     | non-negative formatted monetary quantity. Set to 1 if a space separates the symbol from the              |
| 23336 XSI | value; and set to 2 if a space separates the symbol and the sign string, if adjacent.                    |
| 23337     | <b>char n_cs_precedes</b>                                                                                |
| 23338     | Set to 1 if the <b>currency_symbol</b> or <b>int_curr_symbol</b> precedes the value for a negative       |
| 23339     | formatted monetary quantity. Set to 0 if the symbol succeeds the value.                                  |
| 23340     | <b>char n_sep_by_space</b>                                                                               |
| 23341     | Set to 0 if no space separates the <b>currency_symbol</b> or <b>int_curr_symbol</b> from the value for a |
| 23342     | negative formatted monetary quantity. Set to 1 if a space separates the symbol from the                  |
| 23343 XSI | value; and set to 2 if a space separates the symbol and the sign string, if adjacent.                    |
| 23344     | <b>char p_sign_posn</b>                                                                                  |
| 23345     | Set to a value indicating the positioning of the <b>positive_sign</b> for a non-negative formatted       |
| 23346     | monetary quantity.                                                                                       |
| 23347     | <b>char n_sign_posn</b>                                                                                  |
| 23348     | Set to a value indicating the positioning of the <b>negative_sign</b> for a negative formatted           |
| 23349     | monetary quantity.                                                                                       |
| 23350     | <b>char int_p_cs_precedes</b>                                                                            |
| 23351     | Set to 1 or 0 if the <b>int_currency_symbol</b> respectively precedes or succeeds the value for a        |
| 23352     | non-negative internationally formatted monetary quantity.                                                |
| 23353     | <b>char int_n_cs_precedes</b>                                                                            |
| 23354     | Set to 1 or 0 if the <b>int_currency_symbol</b> respectively precedes or succeeds the value for a        |
| 23355     | negative internationally formatted monetary quantity.                                                    |
| 23356     | <b>char int_p_sep_by_space</b>                                                                           |
| 23357     | Set to a value indicating the separation of the <b>int_currency_symbol</b> , the sign string, and the    |
| 23358     | value for a non-negative internationally formatted monetary quantity.                                    |
| 23359     | <b>char int_n_sep_by_space</b>                                                                           |
| 23360     | Set to a value indicating the separation of the <b>int_currency_symbol</b> , the sign string, and the    |
| 23361     | value for a negative internationally formatted monetary quantity.                                        |
| 23362     | <b>char int_p_sign_posn</b>                                                                              |
| 23363     | Set to a value indicating the positioning of the <b>positive_sign</b> for a non-negative                 |
| 23364     | internationally formatted monetary quantity.                                                             |
| 23365     | <b>char int_n_sign_posn</b>                                                                              |
| 23366     | Set to a value indicating the positioning of the <b>negative_sign</b> for a negative internationally     |

- 23367 formatted monetary quantity.
- 23368 The elements of **grouping** and **mon\_grouping** are interpreted according to the following:
- 23369 {CHAR\_MAX} No further grouping is to be performed.
- 23370 0 The previous element is to be repeatedly used for the remainder of the digits.
- 23371 *other* The integer value is the number of digits that comprise the current group. The  
23372 next element is examined to determine the size of the next group of digits  
23373 before the current group.
- 23374 The values of **p\_sep\_by\_space**, **n\_sep\_by\_space**, **int\_p\_sep\_by\_space**, and **int\_n\_sep\_by\_space**  
23375 are interpreted according to the following:
- 23376 0 No space separates the currency symbol and value.
- 23377 1 If the currency symbol and sign string are adjacent, a space separates them from the value;  
23378 otherwise, a space separates the currency symbol from the value.
- 23379 2 If the currency symbol and sign string are adjacent, a space separates them; otherwise, a  
23380 space separates the sign string from the value.
- 23381 The values of **p\_sign\_posn**, **n\_sign\_posn**, **int\_p\_sign\_posn**, and **int\_n\_sign\_posn** are  
23382 interpreted according to the following:
- 23383 0 Parentheses surround the quantity and **currency\_symbol** or **int\_curr\_symbol**.
- 23384 1 The sign string precedes the quantity and **currency\_symbol** or **int\_curr\_symbol**.
- 23385 2 The sign string succeeds the quantity and **currency\_symbol** or **int\_curr\_symbol**.
- 23386 3 The sign string immediately precedes the **currency\_symbol** or **int\_curr\_symbol**.
- 23387 4 The sign string immediately succeeds the **currency\_symbol** or **int\_curr\_symbol**.
- 23388 The implementation shall behave as if no function in this volume of IEEE Std. 1003.1-200x calls  
23389 *localeconv()*.
- 23390 cx The *localeconv()* function need not be reentrant. A function that is not required to be reentrant is  
23391 not required to be thread-safe.

#### 23392 RETURN VALUE

- 23393 The *localeconv()* function shall return a pointer to the filled-in object. The application shall not  
23394 modify the structure pointed to by the return value which may be overwritten by a subsequent  
23395 call to *localeconv()*. In addition, calls to *setlocale()* with the categories *LC\_ALL*, *LC\_MONETARY*,  
23396 or *LC\_NUMERIC* may overwrite the contents of the structure.

#### 23397 ERRORS

- 23398 No errors are defined.

#### 23399 EXAMPLES

- 23400 None.

#### 23401 APPLICATION USAGE

- 23402 The following table illustrates the rules which may be used by four countries to format monetary  
23403 quantities.

23404  
23405  
23406  
23407  
23408  
23409

| Country     | Positive Format | Negative Format | International Format |
|-------------|-----------------|-----------------|----------------------|
| Italy       | L.1.230         | -L.1.230        | ITL.1.230            |
| Netherlands | F 1.234,56      | F -1.234,56     | NLG 1.234,56         |
| Norway      | kr1.234,56      | kr1.234,56-     | NOK 1.234,56         |
| Switzerland | SFr.s.1,234.56  | SFr.s.1,234.56C | CHF 1,234.56         |

23410  
23411

For these four countries, the respective values for the monetary members of the structure returned by *localeconv()* are:

23412  
23413  
23414  
23415  
23416  
23417  
23418  
23419  
23420  
23421  
23422  
23423  
23424  
23425  
23426  
23427  
23428  
23429  
23430  
23431  
23432  
23433

|                           | Italy   | Netherlands | Norway | Switzerland |
|---------------------------|---------|-------------|--------|-------------|
| <b>int_curr_symbol</b>    | "ITL. " | "NLG "      | "NOK " | "CHF "      |
| <b>currency_symbol</b>    | "L. "   | "F "        | "kr "  | "SFr.s. "   |
| <b>mon_decimal_point</b>  | " "     | ","         | ","    | ."          |
| <b>mon_thousands_sep</b>  | ". "    | ". "        | ". "   | ". "        |
| <b>mon_grouping</b>       | "\3 "   | "\3 "       | "\3 "  | "\3 "       |
| <b>positive_sign</b>      | " "     | " "         | " "    | " "         |
| <b>negative_sign</b>      | "- "    | "- "        | "- "   | "C "        |
| <b>int_frac_digits</b>    | 0       | 2           | 2      | 2           |
| <b>frac_digits</b>        | 0       | 2           | 2      | 2           |
| <b>p_cs_precedes</b>      | 1       | 1           | 1      | 1           |
| <b>p_sep_by_space</b>     | 0       | 1           | 0      | 0           |
| <b>n_cs_precedes</b>      | 1       | 1           | 1      | 1           |
| <b>n_sep_by_space</b>     | 0       | 1           | 0      | 0           |
| <b>p_sign_posn</b>        | 1       | 1           | 1      | 1           |
| <b>n_sign_posn</b>        | 1       | 4           | 2      | 2           |
| <b>int_p_cs_precedes</b>  | 1       | 1           | 1      | 1           |
| <b>int_n_cs_precedes</b>  | 1       | 1           | 1      | 1           |
| <b>int_p_sep_by_space</b> | 0       | 0           | 0      | 0           |
| <b>int_n_sep_by_space</b> | 0       | 0           | 0      | 0           |
| <b>int_p_sign_posn</b>    | 1       | 1           | 1      | 1           |
| <b>int_n_sign_posn</b>    | 1       | 4           | 4      | 2           |

23434 **RATIONALE**  
23435           None.

23436 **FUTURE DIRECTIONS**  
23437           None.

23438 **SEE ALSO**  
23439           *isalpha()*, *isascii()*, *nl\_langinfo()*, *printf()*, *scanf()*, *setlocale()*, *strcat()*, *strchr()*, *strcmp()*, *strcoll()*,  
23440           *strcpy()*, *strftime()*, *strlen()*, *strpbrk()*, *strspn()*, *strtok()*, *strxfrm()*, *strtod()*, the Base Definitions  
23441           volume of IEEE Std. 1003.1-200x, <langinfo.h>, <locale.h>

23442 **CHANGE HISTORY**  
23443           First released in Issue 4. Derived from the ANSI C standard.

23444 **Issue 6**  
23445           A note indicating that this function need not be reentrant is added to the DESCRIPTION.  
23446           The RETURN VALUE section is rewritten to avoid use of the term ‘‘must’’.  
23447           This reference page is updated for alignment with the ISO/IEC 9899: 1999 standard.

23448 **NAME**

23449 localtime, localtime\_r — convert a time value to a broken-down local time

23450 **SYNOPSIS**

23451 #include &lt;time.h&gt;

23452 struct tm \*localtime(const time\_t \*timer);

23453 TSF struct tm \*localtime\_r(const time\_t \*restrict timer,

23454 struct tm \*restrict result);

23455

23456 **DESCRIPTION**

23457 CX For *localtime()*: The functionality described on this reference page is aligned with the ISO C standard. Any conflict between the requirements described here and the ISO C standard is unintentional. This volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

23460 The *localtime()* function shall convert the time in seconds since the Epoch pointed to by *timer* into a broken-down time, expressed as a local time. The function corrects for the timezone and any seasonal time adjustments. Local timezone information is used as though *localtime()* calls *tzset()*.

23464 CX The *localtime()* function need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

23466 TSF The *localtime\_r()* function shall convert the time in seconds since the Epoch pointed to by *timer* into a broken-down time stored in the structure to which *result* points. The *localtime\_r()* function shall also return a pointer to that same structure.

23469 Unlike *localtime()*, the reentrant version is not required to set *tzname*.

23470 **RETURN VALUE**23471 The *localtime()* function shall return a pointer to the broken-down time structure.

23472 TSF Upon successful completion, *localtime\_r()* shall return a pointer to the structure pointed to by the argument *result*.

23474 **ERRORS**

23475 No errors are defined.

23476 **EXAMPLES**23477 **Getting the Local Date and Time**

23478 The following example uses the *time()* function to calculate the time elapsed, in seconds, since January 1, 1970 0:00 UTC (the Epoch), *localtime()* to convert that value to a broken-down time, and *asctime()* to convert the broken-down time values into a printable string.

23481 #include &lt;stdio.h&gt;

23482 #include &lt;time.h&gt;

23483 main()

23484 {

23485 time\_t result;

23486 result = time(NULL);

23487 printf("%s%ld secs since the Epoch\n",

23488 asctime(localtime(&amp;result)),

23489 (long)result);

23490 return(0);

23491 }

23492 This example writes the current time to *stdout* in a form like this:

```
23493 Wed Jun 26 10:32:15 1996
23494 835810335 secs since the Epoch
```

### 23495 **Getting the Modification Time for a File**

23496 The following example gets the modification time for a file. The *localtime()* function converts the **time\_t** value of the last modification date, obtained by a previous call to *stat()*, into a **tm** structure that contains the year, month, day, and so on.

```
23499 #include <time.h>
23500 ...
23501 struct stat statbuf;
23502 ...
23503 tm = localtime(&statbuf.st_mtime);
23504 ...
```

### 23505 **Timing an Event**

23506 The following example gets the current time, converts it to a string using *localtime()* and *asctime()*, and prints it to standard output using *fputs()*. It then prints the number of minutes to an event being timed.

```
23509 #include <time.h>
23510 #include <stdio.h>
23511 ...
23512 time_t now;
23513 int minutes_to_event;
23514 ...
23515 time(&now);
23516 printf("The time is ");
23517 fputs(asctime(localtime(&now)), stdout);
23518 printf("There are still %d minutes to the event.\n",
23519 minutes_to_event);
23520 ...
```

### 23521 **APPLICATION USAGE**

23522 The *asctime()*, *ctime()*, *getdate()*, *gettimeofday()*, *gmtime()*, and *localtime()* functions return values in one of two static objects: a broken-down time structure and an array of **char**. Execution of any of the functions may overwrite the information returned in either of these objects by any of the other functions.

23526 The *localtime\_r()* function is thread-safe and shall return values in a user-supplied buffer instead of possibly using a static data area that may be overwritten by each call.

### 23528 **RATIONALE**

23529 None.

### 23530 **FUTURE DIRECTIONS**

23531 None.

### 23532 **SEE ALSO**

23533 *asctime()*, *clock()*, *ctime()*, *difftime()*, *getdate()*, *gettimeofday()*, *gmtime()*, *mktime()*, *strftime()*,  
23534 *strptime()*, *time()*, *utime()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

23535 **CHANGE HISTORY**

23536 First released in Issue 1. Derived from Issue 1 of the SVID.

23537 **Issue 4**

23538 The APPLICATION USAGE section is expanded to provide a more complete description of how  
23539 static areas are used by the *\*time()* functions.

23540 The following change is incorporated for alignment with the ISO C standard:

- 23541 • The *timer* argument is now a type **const time\_t**.

23542 **Issue 5**

23543 A note indicating that the *localtime()* function need not be reentrant is added to the  
23544 DESCRIPTION.

23545 The *localtime\_r()* function is included for alignment with the POSIX Threads Extension.

23546 **Issue 6**

23547 The *localtime\_r()* function is marked as part of the Thread-Safe Functions option.

23548 Extensions beyond the ISO C standard are now marked.

23549 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
23550 its avoidance of possibly using a static data area.

23551 The **restrict** keyword is added to the *localtime\_r()* prototype for alignment with the  
23552 ISO/IEC 9899:1999 standard.

## 23553 NAME

23554 lockf — record locking on files

## 23555 SYNOPSIS

23556 XSI #include &lt;unistd.h&gt;

23557 int lockf(int *fildes*, int *function*, off\_t *size*);

23558

## 23559 DESCRIPTION

23560 The *lockf()* function allows sections of a file to be locked with advisory-mode locks. Calls to  
 23561 *lockf()* from other threads which attempt to lock the locked file section shall either return an  
 23562 error value or block until the section becomes unlocked. All the locks for a process are removed  
 23563 when the process terminates. Record locking with *lockf()* is supported for regular files and may  
 23564 be supported for other files.

23565 The *fildes* argument is an open file descriptor. The application shall ensure that the file descriptor  
 23566 has been opened with write-only permission (O\_WRONLY) or with read/write permission  
 23567 (O\_RDWR) to establish a lock with this function.

23568 The *function* argument is a control value which specifies the action to be taken. The permissible  
 23569 values for *function* are defined in <unistd.h> as follows:

23570

23571

| Function | Description                                  |
|----------|----------------------------------------------|
| F_ULOCK  | Unlock locked sections.                      |
| F_LOCK   | Lock a section for exclusive use.            |
| F_TLOCK  | Test and lock a section for exclusive use.   |
| F_TEST   | Test a section for locks by other processes. |

23572

23573

23574

23575

23576 F\_TEST detects if a lock by another process is present on the specified section.

23577 F\_LOCK and F\_TLOCK both lock a section of a file if the section is available.

23578 F\_ULOCK removes locks from a section of the file.

23579 The *size* argument is the number of contiguous bytes to be locked or unlocked. The section to be  
 23580 locked or unlocked starts at the current offset in the file and extends forward for a positive *size*  
 23581 or backward for a negative *size* (the preceding bytes up to but not including the current offset).  
 23582 If *size* is 0, the section from the current offset through the largest possible file offset is locked  
 23583 (that is, from the current offset through the present or any future end-of-file). An area need not  
 23584 be allocated to the file to be locked because locks may exist past the end-of-file.

23585 The sections locked with F\_LOCK or F\_TLOCK may, in whole or in part, contain or be contained  
 23586 by a previously locked section for the same process. When this occurs, or if adjacent locked  
 23587 sections would occur, the sections are combined into a single locked section. If the request  
 23588 would cause the number of locks to exceed a system-imposed limit, the request shall fail.

23589 F\_LOCK and F\_TLOCK requests differ only by the action taken if the section is not available.  
 23590 F\_LOCK blocks the calling thread until the section is available. F\_TLOCK makes the function  
 23591 fail if the section is already locked by another process.

23592 File locks are released on first close by the locking process of any file descriptor for the file.

23593 F\_ULOCK requests may release (wholly or in part) one or more locked sections controlled by the  
 23594 process. Locked sections shall be unlocked starting at the current file offset through *size* bytes or  
 23595 to the end-of-file if *size* is (off\_t)0. When all of a locked section is not released (that is, when the  
 23596 beginning or end of the area to be unlocked falls within a locked section), the remaining portions  
 23597 of that section are still locked by the process. Releasing the center portion of a locked section

23598 shall cause the remaining locked beginning and end portions to become two separate locked  
 23599 sections. If the request would cause the number of locks in the system to exceed a system-  
 23600 imposed limit, the request shall fail.

23601 A potential for deadlock occurs if the threads of a process controlling a locked section are  
 23602 blocked by accessing another process' locked section. If the system detects that deadlock would  
 23603 occur, *lockf()* shall fail with an [EDEADLK] error.

23604 The interaction between *fcntl()* and *lockf()* locks is unspecified.

23605 Blocking on a section is interrupted by any signal.

23606 An F\_ULOCK request in which *size* is non-zero and the offset of the last byte of the requested  
 23607 section is the maximum value for an object of type *off\_t*, when the process has an existing lock  
 23608 in which *size* is 0 and which includes the last byte of the requested section, shall be treated as a  
 23609 request to unlock from the start of the requested section with a size equal to 0. Otherwise, an  
 23610 F\_ULOCK request shall attempt to unlock only the requested section.

23611 Attempting to lock a section of a file that is associated with a buffered stream produces  
 23612 unspecified results.

#### 23613 RETURN VALUE

23614 Upon successful completion, *lockf()* shall return 0. Otherwise, it shall return -1, set *errno* to  
 23615 indicate an error, and existing locks shall not be changed.

#### 23616 ERRORS

23617 The *lockf()* function shall fail if:

23618 [EBADF] The *fildev* argument is not a valid open file descriptor; or *function* is F\_LOCK  
 23619 or F\_TLOCK and *fildev* is not a valid file descriptor open for writing.

23620 [EACCES] or [EAGAIN]

23621 The *function* argument is F\_TLOCK or F\_TEST and the section is already  
 23622 locked by another process.

23623 [EDEADLK] The *function* argument is F\_LOCK and a deadlock is detected.

23624 [EINTR] A signal was caught during execution of the function.

23625 [EINVAL] The *function* argument is not one of F\_LOCK, F\_TLOCK, F\_TEST, or  
 23626 F\_ULOCK; or *size* plus the current file offset is less than 0.

23627 [EOVERFLOW] The offset of the first, or if *size* is not 0 then the last, byte in the requested  
 23628 section cannot be represented correctly in an object of type *off\_t*.

23629 The *lockf()* function may fail if:

23630 [EAGAIN] The *function* argument is F\_LOCK or F\_TLOCK and the file is mapped with  
 23631 *mmap()*.

23632 [EDEADLK] or [ENOLCK]

23633 The *function* argument is F\_LOCK, F\_TLOCK, or F\_ULOCK, and the request  
 23634 would cause the number of locks to exceed a system-imposed limit.

23635 [EOPNOTSUPP] or [EINVAL]

23636 The implementation does not support the locking of files of the type indicated  
 23637 by the *fildev* argument.

23638 **EXAMPLES**23639 **Locking a Portion of a File**

23640 In the following example, a file named `/home/cnd/mod1` is being modified. Other processes that  
23641 use locking are prevented from changing it during this process. Only the first 10,000 bytes are  
23642 locked, and the lock call fails if another process has any part of this area locked already.

```
23643 #include <fcntl.h>
23644 #include <unistd.h>
23645 int fildes;
23646 int status;
23647 ...
23648 fildes = open("/home/cnd/mod1", O_RDWR);
23649 status = lockf(fildes, F_TLOCK, (off_t)10000);
```

23650 **APPLICATION USAGE**

23651 Record-locking should not be used in combination with the `fopen()`, `fread()`, `fwrite()`, and other  
23652 `stdio` functions. Instead, the more primitive, non-buffered functions (such as `open()`) should be  
23653 used. Unexpected results may occur in processes that do buffering in the user address space. The  
23654 process may later read/write data which is/was locked. The `stdio` functions are the most  
23655 common source of unexpected buffering.

23656 The `alarm()` function may be used to provide a timeout facility in applications requiring it.

23657 **RATIONALE**

23658 None.

23659 **FUTURE DIRECTIONS**

23660 None.

23661 **SEE ALSO**

23662 `alarm()`, `chmod()`, `close()`, `creat()`, `fcntl()`, `fopen()`, `mmap()`, `open()`, `read()`, `write()`, the Base  
23663 Definitions volume of IEEE Std. 1003.1-200x, `<unistd.h>`

23664 **CHANGE HISTORY**

23665 First released in Issue 4, Version 2.

23666 **Issue 5**

23667 Moved from X/OPEN UNIX extension to BASE.

23668 Large File Summit extensions are added. In particular, the description of [EINVAL] is clarified  
23669 and moved from optional to mandatory status.

23670 A note is added to the DESCRIPTION indicating the effects of attempting to lock a section of a  
23671 file that is associated with a buffered stream.

23672 **Issue 6**

23673 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

23674 **NAME**

23675 log, logf, logl — natural logarithm function

23676 **SYNOPSIS**

23677 #include &lt;math.h&gt;

23678 double log(double x);

23679 float logf(float x);

23680 long double logl(long double x);

23681 **DESCRIPTION**

23682 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 23683 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23684 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

23685 These functions shall compute the natural logarithm of  $x$ ,  $\log_e(x)$ . The application shall ensure  
 23686 that the value of  $x$  is positive.

23687 An application wishing to check for error situations should set *errno* to 0 before calling *log()*. If  
 23688 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

23689 **RETURN VALUE**23690 Upon successful completion, these functions shall return the natural logarithm of  $x$ .23691 XSI If  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].23692 XSI If  $x$  is less than 0, -HUGE\_VAL or NaN shall be returned, and *errno* shall be set to [EDOM].23693 If  $x$  is 0, -HUGE\_VAL shall be returned and *errno* may be set to [ERANGE].23694 **ERRORS**

23695 These functions shall fail if:

23696 [EDOM] The value of  $x$  is negative.

23697 These functions may fail if:

23698 XSI [EDOM] The value of  $x$  is NaN.23699 [ERANGE] The value of  $x$  is 0.

23700 XSI No other errors shall occur.

23701 **EXAMPLES**

23702 None.

23703 **APPLICATION USAGE**

23704 None.

23705 **RATIONALE**

23706 None.

23707 **FUTURE DIRECTIONS**

23708 None.

23709 **SEE ALSO**23710 *exp()*, *isnan()*, *log10()*, *log1p()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>23711 **CHANGE HISTORY**

23712 First released in Issue 1. Derived from Issue 1 of the SVID.

23713 **Issue 4**

23714 References to *matherr()* are removed.

23715 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
23716 ISO C standard and to rationalize error handling in the mathematics functions.

23717 The return value specified for [EDOM] is marked as an extension.

23718 **Issue 5**

23719 The DESCRIPTION is updated to indicate how an application should check for an error. This  
23720 text was previously published in the APPLICATION USAGE section.

23721 **Issue 6**

23722 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

23723 The *logf()* and *logl()* functions are added for alignment with the ISO/IEC 9899:1999 standard. |

23724 **NAME**

23725 log10, log10f, log10l — base 10 logarithm function

23726 **SYNOPSIS**

23727 #include &lt;math.h&gt;

23728 double log10(double x);

23729 float log10f(float x);

23730 long double log10l(long double x);

23731 **DESCRIPTION**

23732 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 23733 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23734 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

23735 These functions shall compute the base 10 logarithm of  $x$ ,  $\log_{10}(x)$ . The application shall ensure  
 23736 that the value of  $x$  is positive.

23737 An application wishing to check for error situations should set *errno* to 0 before calling *log10()*.  
 23738 If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

23739 **RETURN VALUE**23740 Upon successful completion, these functions shall return the base 10 logarithm of  $x$ .23741 **XSI** If  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].23742 **XSI** If  $x$  is less than 0, -HUGE\_VAL or NaN shall be returned, and *errno* shall be set to [EDOM].23743 If  $x$  is 0, -HUGE\_VAL shall be returned and *errno* may be set to [ERANGE].23744 **ERRORS**

23745 These functions shall fail if:

23746 [EDOM] The value of  $x$  is negative.

23747 These functions may fail if:

23748 **XSI** [EDOM] The value of  $x$  is NaN.23749 [ERANGE] The value of  $x$  is 0.23750 **XSI** No other errors shall occur.23751 **EXAMPLES**

23752 None.

23753 **APPLICATION USAGE**

23754 None.

23755 **RATIONALE**

23756 None.

23757 **FUTURE DIRECTIONS**

23758 None.

23759 **SEE ALSO**23760 *isnan()*, *log()*, *pow()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>23761 **CHANGE HISTORY**

23762 First released in Issue 1. Derived from Issue 1 of the SVID.

23763 **Issue 4**

23764       References to *matherr()* are removed.

23765       The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
23766       ISO C standard and to rationalize error handling in the mathematics functions.

23767       The return value specified for [EDOM] is marked as an extension.

23768 **Issue 5**

23769       The DESCRIPTION is updated to indicate how an application should check for an error. This  
23770       text was previously published in the APPLICATION USAGE section.

23771 **Issue 6**

23772       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

23773       The *log10f()* and *log10l()* functions are added for alignment with the ISO/IEC 9899:1999  
23774       standard.

23775 **NAME**

23776 log1p, log1pf, log1pl — compute a natural logarithm

23777 **SYNOPSIS**

23778 #include &lt;math.h&gt;

23779 double log1p(double x);

23780 float log1pf(float x);

23781 long double log1pl(long double x);

23782 **DESCRIPTION**23783 These functions shall compute  $\log_e(1.0 + x)$ . The application shall ensure that the value of  $x$  is  
23784 greater than  $-1.0$ .23785 **RETURN VALUE**23786 Upon successful completion, these functions shall return the natural logarithm of  $1.0 + x$ .23787 If  $x$  is NaN, *log1p()* shall return NaN and may set *errno* to [EDOM].23788 If  $x$  is less than  $-1.0$ , *log1p()* shall return  $-\text{HUGE\_VAL}$  or NaN and set *errno* to [EDOM].23789 If  $x$  is  $-1.0$ , *log1p()* shall return  $-\text{HUGE\_VAL}$  and may set *errno* to [ERANGE].23790 **ERRORS**

23791 These functions shall fail if:

23792 [EDOM] The value of  $x$  is less than  $-1.0$ .23793 These functions may fail and set *errno* to:23794 [EDOM] The value of  $x$  is NaN.23795 [ERANGE] The value of  $x$  is  $-1.0$ .

23796 No other errors shall occur.

23797 **EXAMPLES**

23798 None.

23799 **APPLICATION USAGE**

23800 None.

23801 **RATIONALE**

23802 None.

23803 **FUTURE DIRECTIONS**

23804 None.

23805 **SEE ALSO**23806 *log()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>23807 **CHANGE HISTORY**

23808 First released in Issue 4, Version 2.

23809 **Issue 5**

23810 Moved from X/OPEN UNIX extension to BASE.

23811 **Issue 6**

23812 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

23813 The *log1pf()* and *log1pl()* functions are added for alignment with the ISO/IEC 9899:1999  
23814 standard.

23815 **NAME**

23816 log2, log2f, log2l — compute base 2 logarithm functions

23817 **SYNOPSIS**

23818 #include <math.h>

23819 double log2(double x);

23820 float log2f(float x);

23821 long double log2l(long double x);

23822 **DESCRIPTION**

23823 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
23824 conflict between the requirements described here and the ISO C standard is unintentional. This  
23825 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

23826 These functions shall compute the base 2 logarithm of  $x$ ,  $\log_2(x)$ . The application shall ensure  
23827 that the value of  $x$  is positive.

23828 An application wishing to check for error situations should set *errno* to 0 before calling these  
23829 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

23830 **RETURN VALUE**

23831 Upon successful completion, these functions shall return the base 2 logarithm of  $x$ .

23832 If  $x$  is NaN, these functions shall return NaN and may set *errno* to [EDOM].

23833 If  $x$  is less than 0, these functions shall return  $-\text{HUGE\_VAL}$  or NaN and set *errno* to [EDOM].

23834 If  $x$  is 0, these functions shall return  $-\text{HUGE\_VAL}$  and may set *errno* to [ERANGE].

23835 **ERRORS**

23836 These functions shall fail if:

23837 [EDOM] The value of  $x$  is less than 0.

23838 These functions may fail if:

23839 [EDOM] The value of  $x$  is NaN.

23840 [ERANGE] The value of  $x$  is 0.

23841 No other errors shall occur.

23842 **EXAMPLES**

23843 None.

23844 **APPLICATION USAGE**

23845 None.

23846 **RATIONALE**

23847 None.

23848 **FUTURE DIRECTIONS**

23849 None.

23850 **SEE ALSO**

23851 *log()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

23852 **CHANGE HISTORY**

23853 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

23854 **NAME**

23855 logb, logbf, logbl — radix-independent exponent

23856 **SYNOPSIS**

23857 XSI #include &lt;math.h&gt;

23858 double logb(double x);

23859 float logbf(float x);

23860 long double logbl(long double x);

23861

23862 **DESCRIPTION**

23863 These functions shall compute the exponent of  $x$ , which is the integral part of  $\log_r |x|$ , as a  
 23864 signed floating point value, for non-zero  $x$ , where  $r$  is the radix of the machine's floating-point  
 23865 arithmetic, which is the value of FLT\_RADIX defined in the <float.h> header.

23866 **RETURN VALUE**23867 Upon successful completion, these functions shall return the exponent of  $x$ .23868 If  $x$  is 0.0, *logb()* shall return `-HUGE_VAL` and set *errno* to [EDOM].23869 If  $x$  is  $\pm\text{Inf}$ , *logb()* shall return `+Inf`.23870 If  $x$  is NaN, *logb()* shall return NaN and may set *errno* to [EDOM].23871 **ERRORS**

23872 These functions shall fail if:

23873 [EDOM] The  $x$  argument is 0.0.

23874 These functions may fail if:

23875 [EDOM] The  $x$  argument is NaN.23876 **EXAMPLES**

23877 None.

23878 **APPLICATION USAGE**

23879 None.

23880 **RATIONALE**

23881 None.

23882 **FUTURE DIRECTIONS**

23883 None.

23884 **SEE ALSO**23885 *ilogb()*, *scalb()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <float.h> <math.h>23886 **CHANGE HISTORY**

23887 First released in Issue 4, Version 2.

23888 **Issue 5**

23889 Moved from X/OPEN UNIX extension to BASE.

23890 **Issue 6**23891 The *logbf()* and *logbl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

## 23892 NAME

23893 longjmp — non-local goto

## 23894 SYNOPSIS

23895 #include &lt;setjmp.h&gt;

23896 void longjmp(jmp\_buf env, int val);

## 23897 DESCRIPTION

23898 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 23899 conflict between the requirements described here and the ISO C standard is unintentional. This  
 23900 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

23901 The *longjmp()* function shall restore the environment saved by the most recent invocation of  
 23902 *setjmp()* in the same thread, with the corresponding **jmp\_buf** argument. If there is no such  
 23903 invocation, or if the function containing the invocation of the *setjmp()* macro has terminated  
 23904 execution in the interim, or if the invocation of *setjmp()* was within the scope of an identifier  
 23905 with variably modified type and execution has left that scope in the interim, the behavior is  
 23906 cx undefined. It is unspecified whether *longjmp()* restores the signal mask, leaves the signal mask  
 23907 unchanged, or restores it to its value at the time *setjmp()* was called.

23908 All accessible objects have values as of the time *longjmp()* was called, and all other components  
 23909 of the abstract machine have state (for example, floating-point status flags and open files),  
 23910 except that the values of objects of automatic storage duration are indeterminate if they meet all  
 23911 the following conditions:

- 23912 • They are local to the function containing the corresponding *setjmp()* invocation.
- 23913 • They do not have volatile-qualified type.
- 23914 • They are changed between the *setjmp()* invocation and *longjmp()* call.

23915 As it bypasses the usual function call and return mechanisms, *longjmp()* shall execute correctly  
 23916 in contexts of interrupts, signals, and any of their associated functions. However, if *longjmp()* is  
 23917 invoked from a nested signal handler (that is, from a function invoked as a result of a signal  
 23918 raised during the handling of another signal), the behavior is undefined.

23919 cx The effect of a call to *longjmp()* where initialization of the **jmp\_buf** structure was not performed  
 23920 in the calling thread is undefined.

## 23921 RETURN VALUE

23922 After *longjmp()* is completed, program execution continues as if the corresponding invocation of  
 23923 *setjmp()* had just returned the value specified by *val*. The *longjmp()* function shall not cause  
 23924 *setjmp()* to return 0; if *val* is 0, *setjmp()* shall return 1.

## 23925 ERRORS

23926 No errors are defined.

## 23927 EXAMPLES

23928 None.

## 23929 APPLICATION USAGE

23930 Applications whose behavior depends on the value of the signal mask should not use *longjmp()*  
 23931 and *setjmp()*, since their effect on the signal mask is unspecified, but should instead use the  
 23932 *siglongjmp()* and *sigsetjmp()* functions (which can save and restore the signal mask under  
 23933 application control).

23934 **RATIONALE**

23935 None.

23936 **FUTURE DIRECTIONS**

23937 None.

23938 **SEE ALSO**

23939 *setjmp()*, *sigaction()*, *siglongjmp()*, *sigsetjmp()*, the Base Definitions volume of  
23940 IEEE Std. 1003.1-200x, <**setjmp.h**>

23941 **CHANGE HISTORY**

23942 First released in Issue 1. Derived from Issue 1 of the SVID.

23943 **Issue 4**

23944 The APPLICATION USAGE section is deleted.

23945 The following change is incorporated for alignment with the ISO C standard:

23946

- Mention of volatile-qualified types is added to the DESCRIPTION.

23947 **Issue 4, Version 2**

23948 The DESCRIPTION is updated for X/OPEN UNIX conformance and discusses valid possibilities  
23949 for the resulting state of the signal mask.

23950 **Issue 5**

23951 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

23952 **Issue 6**

23953 Extensions beyond the ISO C standard are now marked.

23954 The following new requirements on POSIX implementations derive from alignment with the  
23955 Single UNIX Specification:

23956

- The DESCRIPTION now explicitly makes *longjmp()*'s effect on the signal mask unspecified.

23957 The DESCRIPTION is updated for alignment with the ISO/IEC 9899:1999 standard.

23958 **NAME**

23959       lrand48 — generate uniformly distributed pseudo-random non-negative long integers

23960 **SYNOPSIS**

23961 xSI     #include <stdlib.h>

23962       long lrand48(void);

23963

23964 **DESCRIPTION**

23965       Refer to *drand48()*.

23966 **NAME**

23967 lsearch, lfind — linear search and update

23968 **SYNOPSIS**

```
23969 xSI #include <search.h>
23970 void *lsearch(const void *key, void *base, size_t *nel, size_t width,
23971 int (*compar)(const void *, const void *));
23972 void *lfind(const void *key, const void *base, size_t *nel,
23973 size_t width, int (*compar)(const void *, const void *));
23974
```

23975 **DESCRIPTION**

23976 The *lsearch()* function is a linear search routine. It returns a pointer into the table for the  
 23977 matching entry. If the entry does not occur, it is added at the end of the table. The *key* argument  
 23978 points to the entry to be sought in the table. The *base* argument points to the first element in the  
 23979 table. The *width* argument is the size of an element in bytes. The *nel* argument points to an  
 23980 integer containing the current number of elements in the table. The integer to which *nel* points  
 23981 is incremented if the entry is added to the table. The *compar* argument points to a comparison  
 23982 function which the application shall supply (for example, *strcmp()*). It is called with two  
 23983 arguments that point to the elements being compared. The application shall ensure that the  
 23984 function returns 0 if the elements are equal, and non-zero otherwise.

23985 The *lfind()* function is the same as *lsearch()* except that if the entry is not found, it is not added to  
 23986 the table. Instead, a null pointer is returned.

23987 **RETURN VALUE**

23988 If the searched for entry is found, both *lsearch()* and *lfind()* shall return a pointer to it. Otherwise,  
 23989 *lfind()* shall return a null pointer and *lsearch()* shall return a pointer to the newly added element.

23990 Both functions shall return a null pointer in case of error.

23991 **ERRORS**

23992 No errors are defined.

23993 **EXAMPLES**23994 **Storing Strings in a Table**

23995 This fragment reads in less than or equal to TABSIZE strings of length less than or equal to  
 23996 ELSIZE and stores them in a table, eliminating duplicates.

```
23997 #include <stdio.h>
23998 #include <string.h>
23999 #include <search.h>
24000 #define TABSIZE 50
24001 #define ELSIZE 120
24002 ...
24003 char line[ELSIZE], tab[TABSIZE][ELSIZE];
24004 size_t nel = 0;
24005 ...
24006 while (fgets(line, ELSIZE, stdin) != NULL && nel < TABSIZE)
24007 (void) lsearch(line, tab, &nel,
24008 ELSIZE, (int (*)(const void *, const void *)) strcmp);
24009 ...
```

24010 **Finding a Matching Entry**

24011 The following example finds any line that reads "This is a test.".

```
24012 #include <search.h>
24013 #include <string.h>
24014 ...
24015 char line[ELSIZE], tab[TABSIZE][ELSIZE];
24016 size_t nel = 0;
24017 char *findline;
24018 void *entry;

24019 findline = "This is a test.\n";

24020 entry = lfind(findline, tab, &nel, ELSIZE, (
24021 int (*)(const void *, const void *)) strcmp);
```

24022 **APPLICATION USAGE**

24023 The comparison function need not compare every byte, so arbitrary data may be contained in  
24024 the elements in addition to the values being compared.

24025 Undefined results can occur if there is not enough room in the table to add a new item.

24026 **RATIONALE**

24027 None.

24028 **FUTURE DIRECTIONS**

24029 None.

24030 **SEE ALSO**

24031 *hcreate()*, *tsearch()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**search.h**>

24032 **CHANGE HISTORY**

24033 First released in Issue 1. Derived from Issue 1 of the SVID.

24034 **Issue 4**

24035 In the SYNOPSIS section, the type of argument *key* in the declaration of *lsearch()* is changed from  
24036 **void\*** to **const void\***, the type arguments *key* and *base* have been changed from **void\*** to **const**  
24037 **void\*** in the declaration of *lfind()*, and the arguments to *compar* are defined for both functions.

24038 In the EXAMPLES section, the sample code is updated to use ISO C standard syntax.

24039 Warnings about the casting of various arguments are removed from the APPLICATION USAGE  
24040 section, as casting requirements are now clear from the function definitions.

24041 **Issue 6**

24042 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

24043 **NAME**

24044 lseek — move the read/write file offset

24045 **SYNOPSIS**

24046 #include &lt;unistd.h&gt;

24047 off\_t lseek(int *fildev*, off\_t *offset*, int *whence*);24048 **DESCRIPTION**24049 The *lseek()* function shall set the file offset for the open file description associated with the file descriptor *fildev*, as follows:

- 24051 • If *whence* is {SEEK\_SET}, the file offset is set to *offset* bytes.
- 24052 • If *whence* is {SEEK\_CUR}, the file offset is set to its current location plus *offset*.
- 24053 • If *whence* is {SEEK\_END}, the file offset is set to the size of the file plus *offset*.

24054 The symbolic constants {SEEK\_SET}, {SEEK\_CUR}, and {SEEK\_END} are defined in &lt;unistd.h&gt;.

24055 The behavior of *lseek()* on devices which are incapable of seeking is implementation-defined.  
24056 The value of the file offset associated with such a device is undefined.24057 The *lseek()* function shall allow the file offset to be set beyond the end of the existing data in the file. If data is later written at this point, subsequent reads of data in the gap shall return bytes with the value 0 until data is actually written into the gap.24060 The *lseek()* function shall not, by itself, extend the size of a file.24061 SHM If *fildev* refers to a shared memory object, the result of the *lseek()* function is unspecified.24062 TYM If *fildev* refers to a typed memory object, the result of the *lseek()* function is unspecified.24063 **RETURN VALUE**24064 Upon successful completion, the resulting offset, as measured in bytes from the beginning of the file, shall be returned. Otherwise, (off\_t)-1 shall be returned, *errno* shall be set to indicate the error, and the file offset shall remain unchanged.24067 **ERRORS**24068 The *lseek()* function shall fail if:

- |       |             |                                                                                                                                                        |
|-------|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| 24069 | [EBADF]     | The <i>fildev</i> argument is not an open file descriptor.                                                                                             |
| 24070 | [EINVAL]    | The <i>whence</i> argument is not a proper value, or the resulting file offset would be negative for a regular file, block special file, or directory. |
| 24072 | [EOVERFLOW] | The resulting file offset would be a value which cannot be represented correctly in an object of type <b>off_t</b> .                                   |
| 24074 | [ESPIPE]    | The <i>fildev</i> argument is associated with a pipe, FIFO, or socket.                                                                                 |

24075 **EXAMPLES**

24076 None.

24077 **APPLICATION USAGE**

24078 None.

24079 **RATIONALE**24080 The ISO C standard includes the functions *fgetpos()* and *fsetpos()*, which work on very large files by use of a special positioning type.24082 Although *lseek()* may position the file offset beyond the end of the file, this function does not itself extend the size of the file. While the only function in this volume of IEEE Std. 1003.1-200x

- 24084 that may directly extend the size of the file is *write()*, several functions originally derived from  
24085 the ISO C standard, such as *fwrite()*, *fprintf()*, and so on, may do so (by causing calls on *write()*).
- 24086 An invalid file offset that would cause [EINVAL] to be returned may be both implementation-  
24087 defined and device-dependent (for example, memory may have few invalid values). A negative  
24088 file offset may be valid for some devices in some implementations.
- 24089 The POSIX.1-1990 standard did not specifically prohibit *lseek()* from returning a negative offset.  
24090 Therefore, an application was required to clear *errno* prior to the call and check *errno* upon return  
24091 to determine whether a return value of (*off\_t*)-1 is a negative offset or an indication of an error  
24092 condition. The standard developers did not wish to require this action on the part of a portable  
24093 application, and chose to require that *errno* be set to [EINVAL] when the resulting file offset  
24094 would be negative for a regular file, block special file, or directory.
- 24095 **FUTURE DIRECTIONS**
- 24096 None.
- 24097 **SEE ALSO**
- 24098 *open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>
- 24099 **CHANGE HISTORY**
- 24100 First released in Issue 1. Derived from Issue 1 of the SVID.
- 24101 **Issue 4**
- 24102 The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
24103 XSI-conformant systems.
- 24104 The APPLICATION USAGE section is removed, as the ISO POSIX-1 standard now requires that  
24105 *off\_t* be signed.
- 24106 **Issue 5**
- 24107 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension.
- 24108 Large File Summit extensions are added.
- 24109 **Issue 6**
- 24110 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.
- 24111 The following new requirements on POSIX implementations derive from alignment with the  
24112 Single UNIX Specification:
- 24113 • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
24114 required for conforming implementations of previous POSIX specifications, it was not  
24115 required for UNIX applications.
  - 24116 • The [EOVERFLOW] error condition is added. This change is to support large files.
- 24117 An additional [ESPIPE] error condition is added for sockets.
- 24118 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that  
24119 *lseek()* results are unspecified for typed memory objects.

24120 **NAME**

24121 lstat — get symbolic link status

24122 **SYNOPSIS**

24123 #include &lt;sys/stat.h&gt;

24124 int lstat(const char \*restrict *path*, struct stat \*restrict *buf*);24125 **DESCRIPTION**

24126 The *lstat()* function shall have the same effect as *stat()*, except when *path* refers to a symbolic  
 24127 link. In that case *lstat()* shall return information about the link, while *stat()* shall return  
 24128 information about the file the link references.

24129 For symbolic links, the *st\_mode* member shall contain meaningful information when used with  
 24130 the file type macros, and the *st\_size* member shall contain the length of the path name contained  
 24131 in the symbolic link. File mode bits and the contents of the remaining members of the **stat**  
 24132 structure are unspecified. The value returned in the *st\_size* member is the length of the contents  
 24133 of the symbolic link, and does not count any trailing null.

24134 **RETURN VALUE**

24135 Upon successful completion, *lstat()* shall return 0. Otherwise, it shall return -1 and set *errno* to  
 24136 indicate the error.

24137 **ERRORS**24138 The *lstat()* function shall fail if:

24139 [EACCES] A component of the path prefix denies search permission.

24140 [EIO] An error occurred while reading from the file system.

24141 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 24142 argument.

24143 [ENAMETOOLONG]

24144 The length of a path name exceeds {PATH\_MAX} or a path name component  
 24145 is longer than {NAME\_MAX}.

24146 [ENOTDIR] A component of the path prefix is not a directory.

24147 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

24148 [EOVERFLOW] The file size in bytes or the number of blocks allocated to the file or the file  
 24149 serial number cannot be represented correctly in the structure pointed to by  
 24150 *buf*.

24151 The *lstat()* function may fail if:

24152 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 24153 resolution of the *path* argument.

24154 [ENAMETOOLONG]

24155 As a result of encountering a symbolic link in resolution of the *path* argument,  
 24156 the length of the substituted path name string exceeded {PATH\_MAX}.

24157 [EOVERFLOW] One of the members is too large to store into the structure pointed to by the  
 24158 *buf* argument.

24159 **EXAMPLES**24160 **Obtaining Symbolic Link Status Information**

24161 The following example shows how to obtain status information for a symbolic link named  
24162 **/modules/pass1**. The structure variable *buffer* is defined for the **stat** structure. If the *path*  
24163 argument specified the file name for the file pointed to by the symbolic link (**/home/cnd/mod1**),  
24164 the results of calling the function would be the same as those returned by a call to the *stat()*  
24165 function.

```
24166 #include <sys/stat.h>
24167 struct stat buffer;
24168 int status;
24169 ...
24170 status = lstat("/modules/pass1", &buffer);
```

24171 **APPLICATION USAGE**

24172 None.

24173 **RATIONALE**

24174 The *lstat()* function is not required to update the time-related fields if the named file is not a  
24175 symbolic link. While the *st\_uid*, *st\_gid*, *st\_atime*, *st\_mtime*, and *st\_ctime* members of the **stat**  
24176 structure may apply to a symbolic link, they are not required to do so. No functions in  
24177 IEEE Std. 1003.1-200x are required to maintain any of these time fields.

24178 **FUTURE DIRECTIONS**

24179 None.

24180 **SEE ALSO**

24181 *lstat()*, *readlink()*, *stat()*, *symlink()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
24182 **<sys/stat.h>**

24183 **CHANGE HISTORY**

24184 First released in Issue 4, Version 2.

24185 **Issue 5**

24186 Moved from X/OPEN UNIX extension to BASE.

24187 Large File Summit extensions are added.

24188 **Issue 6**

24189 The following changes were made to align with the IEEE P1003.1a draft standard:

- 24190 • This function is now mandatory.
- 24191 • The [ELOOP] optional error condition is added.

24192 The **restrict** keyword is added to the *lstat()* prototype for alignment with the ISO/IEC 9899:1999  
24193 standard.

24194 **NAME**

24195       makecontext, swapcontext — manipulate user contexts

24196 **SYNOPSIS**

24197 XSI       #include &lt;ucontext.h&gt;

24198       void makecontext(ucontext\_t \*ucp, void (\*func)(void),  
24199                       int argc, ...);

24200       int swapcontext(ucontext\_t \*restrict oucp,

24201                       const ucontext\_t \*restrict ucp);

24202

24203 **DESCRIPTION**

24204       The *makecontext()* function shall modify the context specified by *ucp*, which has been initialized  
24205       using *getcontext()*. When this context is resumed using *swapcontext()* or *setcontext()*, program  
24206       execution shall continue by calling *func*, passing it the arguments that follow *argc* in the  
24207       *makecontext()* call.

24208       Before a call is made to *makecontext()*, the application shall ensure that the context being  
24209       modified has a stack allocated for it. The application shall ensure that the value of *argc* matches  
24210       the number of integer arguments passed to *func*; otherwise, the behavior is undefined.

24211       The *uc\_link* member is used to determine the context that shall be resumed when the context  
24212       being modified by *makecontext()* returns. The application shall ensure that the *uc\_link* member is  
24213       initialized prior to the call to *makecontext()*.

24214       The *swapcontext()* function shall save the current context in the context structure pointed to by  
24215       *oucp* and shall set the context to the context structure pointed to by *ucp*.

24216 **RETURN VALUE**

24217       Upon successful completion, *swapcontext()* shall return 0. Otherwise,  $-1$  shall be returned and  
24218       *errno* set to indicate the error.

24219 **ERRORS**24220       The *swapcontext()* function shall fail if:24221       [ENOMEM]       The *ucp* argument does not have enough stack left to complete the operation.24222 **EXAMPLES**

24223       None.

24224 **APPLICATION USAGE**

24225       None.

24226 **RATIONALE**

24227       None.

24228 **FUTURE DIRECTIONS**

24229       None.

24230 **SEE ALSO**

24231       *exit()*, *getcontext()*, *sigaction()*, *sigprocmask()*, the Base Definitions volume of  
24232       IEEE Std. 1003.1-200x, <ucontext.h>

24233 **CHANGE HISTORY**

24234       First released in Issue 4, Version 2.

24235 **Issue 5**

24236 Moved from X/OPEN UNIX extension to BASE.

24237 In the ERRORS section, the description of [ENOMEM] is changed to apply to *swapcontext()* only.

24238 **Issue 6**

24239 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

24240 The **restrict** keyword is added to the *swapcontext()* prototype for alignment with the

24241 ISO/IEC 9899:1999 standard.

24242 **NAME**

24243 malloc — a memory allocator

24244 **SYNOPSIS**

24245 #include &lt;stdlib.h&gt;

24246 void \*malloc(size\_t size);

24247 **DESCRIPTION**

24248 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 24249 conflict between the requirements described here and the ISO C standard is unintentional. This  
 24250 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24251 The *malloc()* function shall allocate unused space for an object whose size in bytes is specified by  
 24252 *size* and whose value is indeterminate.

24253 The order and contiguity of storage allocated by successive calls to *malloc()* is unspecified. The  
 24254 pointer returned if the allocation succeeds shall be suitably aligned so that it may be assigned to  
 24255 a pointer to any type of object and then used to access such an object in the space allocated (until  
 24256 the space is explicitly freed or reallocated). Each such allocation shall yield a pointer to an object  
 24257 disjoint from any other object. The pointer returned points to the start (lowest byte address) of  
 24258 the allocated space. If the space cannot be allocated, a null pointer shall be returned. If the size of  
 24259 the space requested is 0, the behavior is implementation-defined; the value returned shall be  
 24260 either a null pointer or a unique pointer.

24261 **RETURN VALUE**

24262 Upon successful completion with *size* not equal to 0, *malloc()* shall return a pointer to the  
 24263 allocated space. If *size* is 0, either a null pointer or a unique pointer that can be successfully  
 24264 CX passed to *free()* shall be returned. Otherwise, it shall return a null pointer and set *errno* to  
 24265 indicate the error.

24266 **ERRORS**24267 The *malloc()* function shall fail if:

24268 CX [ENOMEM] Insufficient storage space is available.

24269 **EXAMPLES**

24270 None.

24271 **APPLICATION USAGE**

24272 None.

24273 **RATIONALE**

24274 None.

24275 **FUTURE DIRECTIONS**

24276 None.

24277 **SEE ALSO**24278 *calloc()*, *free()*, *realloc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>24279 **CHANGE HISTORY**

24280 First released in Issue 1. Derived from Issue 1 of the SVID.

24281 **Issue 4**24282 The setting of *errno* and the [ENOMEM] error are marked as extensions.

24283 The APPLICATION USAGE section is changed to record that <malloc.h> need no longer be  
 24284 supported on XSI-conformant systems.

- 24285 The following change is incorporated for alignment with the ISO C standard:
- 24286
  - The RETURN VALUE section is updated to indicate what is returned if *size* is 0.
- 24287 **Issue 6**
- 24288 Extensions beyond the ISO C standard are now marked.
- 24289 The following new requirements on POSIX implementations derive from alignment with the  
24290 Single UNIX Specification:
- 24291
  - In the RETURN VALUE section, the requirement to set *errno* to indicate an error is added.
- 24292
  - The [ENOMEM] error condition is added.

24293 **NAME**

24294           mblen — get number of bytes in a character

24295 **SYNOPSIS**

24296           #include &lt;stdlib.h&gt;

24297           int mblen(const char \*s, size\_t n);

24298 **DESCRIPTION**

24299 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 24300 conflict between the requirements described here and the ISO C standard is unintentional. This  
 24301 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24302       If *s* is not a null pointer, *mblen()* determines the number of bytes constituting the character  
 24303 pointed to by *s*. Except that the shift state of *mbtowc()* is not affected, it is equivalent to:

24304       mbtowc((wchar\_t \*)0, s, n);

24305       The implementation shall behave as if no function defined in this volume of  
 24306 IEEE Std. 1003.1-200x calls *mblen()*.

24307       The behavior of this function is affected by the *LC\_CTYPE* category of the current locale. For a  
 24308 state-dependent encoding, this function is placed into its initial state by a call for which its  
 24309 character pointer argument, *s*, is a null pointer. Subsequent calls with *s* as other than a null  
 24310 pointer cause the internal state of the function to be altered as necessary. A call with *s* as a null  
 24311 pointer causes this function to return a non-zero value if encodings have state dependency, and  
 24312 0 otherwise. If the implementation employs special bytes to change the shift state, these bytes do  
 24313 not produce separate wide-character codes, but are grouped with an adjacent character.  
 24314 Changing the *LC\_CTYPE* category causes the shift state of this function to be indeterminate.

24315 **RETURN VALUE**

24316       If *s* is a null pointer, *mblen()* shall return a non-zero or 0 value, if character encodings,  
 24317 respectively, do or do not have state-dependent encodings. If *s* is not a null pointer, *mblen()* shall  
 24318 either return 0 (if *s* points to the null byte), or return the number of bytes that constitute the  
 24319 character (if the next *n* or fewer bytes form a valid character), or return -1 (if they do not form a  
 24320 **CX** valid character) and may set *errno* to indicate the error. In no case shall the value returned be  
 24321 greater than *n* or the value of the {MB\_CUR\_MAX} macro.

24322 **ERRORS**24323       The *mblen()* function may fail if:24324 **XSI**       [EILSEQ]       Invalid character sequence is detected.24325 **EXAMPLES**

24326       None.

24327 **APPLICATION USAGE**

24328       None.

24329 **RATIONALE**

24330       None.

24331 **FUTURE DIRECTIONS**

24332       None.

24333 **SEE ALSO**

24334       *mbtowc()*, *mbstowcs()*, *wctomb()*, *wcstombs()*, the Base Definitions volume of  
 24335 IEEE Std. 1003.1-200x, <stdlib.h>

24336 **CHANGE HISTORY**

24337 First released in Issue 4. Aligned with the ISO C standard.

24338 **NAME**24339 `mbrlen` — get number of bytes in a character (restartable)24340 **SYNOPSIS**24341 `#include <wchar.h>`24342 `size_t mbrlen(const char *restrict s, size_t n,`  
24343 `mbstate_t *restrict ps);`24344 **DESCRIPTION**24345 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
24346 conflict between the requirements described here and the ISO C standard is unintentional. This  
24347 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.24348 If *s* is not a null pointer, `mbrlen()` shall determine the number of bytes constituting the character  
24349 pointed to by *s*. It is equivalent to:24350 `mbstate_t internal;`24351 `mbrtowc(NULL, s, n, ps != NULL ? ps : &internal);`24352 If *ps* is a null pointer, the `mbrlen()` function uses its own internal `mbstate_t` object, which is  
24353 initialized at program start-up to the initial conversion state. Otherwise, the `mbstate_t` object  
24354 pointed to by *ps* is used to completely describe the current conversion state of the associated  
24355 character sequence. The implementation shall behave as if no function defined in this volume of  
24356 IEEE Std. 1003.1-200x calls `mbrlen()`.24357 **XSI** The behavior of this function is affected by the `LC_CTYPE` category of the current locale.24358 **RETURN VALUE**24359 The `mbrlen()` function shall return the first of the following that applies:24360 **0** If the next *n* or fewer bytes complete the character that corresponds to the null  
24361 wide character.24362 **positive** If the next *n* or fewer bytes complete a valid character; the value returned shall  
24363 be the number of bytes that complete the character.24364 **(size\_t)-2** If the next *n* bytes contribute to an incomplete but potentially valid character,  
24365 and all *n* bytes have been processed. When *n* has at least the value of the  
24366 `{MB_CUR_MAX}` macro, this case can only occur if *s* points at a sequence of  
24367 redundant shift sequences (for implementations with state-dependent  
24368 encodings).24369 **(size\_t)-1** If an encoding error occurs, in which case the next *n* or fewer bytes do not  
24370 contribute to a complete and valid character. In this case, `[EILSEQ]` shall be  
24371 stored in `errno` and the conversion state is undefined.24372 **ERRORS**24373 The `mbrlen()` function may fail if:24374 `[EINVAL]` *ps* points to an object that contains an invalid conversion state.24375 `[EILSEQ]` Invalid character sequence is detected.

24376 **EXAMPLES**

24377 None.

24378 **APPLICATION USAGE**

24379 None.

24380 **RATIONALE**

24381 None.

24382 **FUTURE DIRECTIONS**

24383 None.

24384 **SEE ALSO**24385 *mbsinit()*, *mbrtowc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<wchar.h>**24386 **CHANGE HISTORY**

24387 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995

24388 (E).

24389 **Issue 6**24390 The *mbrlen()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

## 24391 NAME

24392 mbrtowc — convert a character to a wide-character code (restartable)

## 24393 SYNOPSIS

24394 #include &lt;wchar.h&gt;

```
24395 size_t mbrtowc(wchar_t *restrict pwc, const char *restrict s,
24396 size_t n, mbstate_t *restrict ps);
```

## 24397 DESCRIPTION

24398 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 24399 conflict between the requirements described here and the ISO C standard is unintentional. This  
 24400 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24401 If *s* is a null pointer, the *mbrtowc()* function shall be equivalent to the call:

```
24402 mbrtowc(NULL, "", 1, ps)
```

24403 In this case, the values of the arguments *pwc* and *n* are ignored.

24404 If *s* is not a null pointer, the *mbrtowc()* function inspects at most *n* bytes beginning at the byte  
 24405 pointed to by *s* to determine the number of bytes needed to complete the next character  
 24406 (including any shift sequences). If the function determines that the next character is completed, it  
 24407 determines the value of the corresponding wide character and then, if *pwc* is not a null pointer,  
 24408 stores that value in the object pointed to by *pwc*. If the corresponding wide character is the null  
 24409 wide character, the resulting state described is the initial conversion state.

24410 If *ps* is a null pointer, the *mbrtowc()* function uses its own internal **mbstate\_t** object, which is  
 24411 initialized at program start-up to the initial conversion state. Otherwise, the **mbstate\_t** object  
 24412 pointed to by *ps* is used to completely describe the current conversion state of the associated  
 24413 character sequence. The implementation shall behave as if no function defined in this volume of  
 24414 IEEE Std. 1003.1-200x calls *mbrtowc()*.

24415 xsi The behavior of this function is affected by the *LC\_CTYPE* category of the current locale.

## 24416 RETURN VALUE

24417 The *mbrtowc()* function shall return the first of the following that applies:

24418 0 If the next *n* or fewer bytes complete the character that corresponds to the null  
 24419 wide character (which is the value stored).

24420 *positive* If the next *n* or fewer bytes complete a valid character (which is the value  
 24421 stored); the value returned shall be the number of bytes that complete the  
 24422 character.

24423 **(size\_t)–2** If the next *n* bytes contribute to an incomplete but potentially valid character,  
 24424 and all *n* bytes have been processed (no value is stored). When *n* has at least  
 24425 the value of the {*MB\_CUR\_MAX*} macro, this case can only occur if *s* points at  
 24426 a sequence of redundant shift sequences (for implementations with state-  
 24427 dependent encodings).

24428 **(size\_t)–1** If an encoding error occurs, in which case the next *n* or fewer bytes do not  
 24429 contribute to a complete and valid character (no value is stored). In this case,  
 24430 [EILSEQ] shall be stored in *errno* and the conversion state is undefined.

## 24431 ERRORS

24432 The *mbrtowc()* function may fail if:

24433 cx [EINVAL] *ps* points to an object that contains an invalid conversion state.

- 24434 [EILSEQ] Invalid character sequence is detected.
- 24435 **EXAMPLES**
- 24436 None.
- 24437 **APPLICATION USAGE**
- 24438 None.
- 24439 **RATIONALE**
- 24440 None.
- 24441 **FUTURE DIRECTIONS**
- 24442 None.
- 24443 **SEE ALSO**
- 24444 *mbstinit()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>
- 24445 **CHANGE HISTORY**
- 24446 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995
- 24447 (E).
- 24448 **Issue 6**
- 24449 The *mbrtowc()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

24450 **NAME**

24451 mbsinit — determine conversion object status

24452 **SYNOPSIS**

24453 #include <wchar.h>

24454 int mbsinit(const mbstate\_t \*ps);

24455 **DESCRIPTION**

24456 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
24457 conflict between the requirements described here and the ISO C standard is unintentional. This  
24458 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24459 If *ps* is not a null pointer, the *mbsinit()* function shall determine whether the object pointed to by  
24460 *ps* describes an initial conversion state.

24461 **RETURN VALUE**

24462 The *mbsinit()* function shall return non-zero if *ps* is a null pointer, or if the pointed-to object  
24463 describes an initial conversion state; otherwise, it shall return zero.

24464 If an **mbstate\_t** object is altered by any of the functions described as “restartable”, and is then  
24465 used with a different character sequence, or in the other conversion direction, or with a different  
24466 *LC\_CTYPE* category setting than on earlier function calls, the behavior is undefined.

24467 **ERRORS**

24468 No errors are defined.

24469 **EXAMPLES**

24470 None.

24471 **APPLICATION USAGE**

24472 The **mbstate\_t** object is used to describe the current conversion state from a particular character  
24473 sequence to a wide-character sequence (or *vice versa*) under the rules of a particular setting of the  
24474 *LC\_CTYPE* category of the current locale.

24475 The initial conversion state corresponds, for a conversion in either direction, to the beginning of  
24476 a new character sequence in the initial shift state. A zero valued **mbstate\_t** object is at least one  
24477 way to describe an initial conversion state. A zero valued **mbstate\_t** object can be used to initiate  
24478 conversion involving any character sequence, in any *LC\_CTYPE* category setting.

24479 **RATIONALE**

24480 None.

24481 **FUTURE DIRECTIONS**

24482 None.

24483 **SEE ALSO**

24484 *mbrlen()*, *mbrtowc()*, *wcrtomb()*, *mbsrtowcs()*, *wcsrtombs()*, the Base Definitions volume of  
24485 IEEE Std. 1003.1-200x, <**wchar.h**>

24486 **CHANGE HISTORY**

24487 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
24488 (E).

24489 **NAME**

24490 mbsrtowcs — convert a character string to a wide-character string (restartable)

24491 **SYNOPSIS**

24492 #include &lt;wchar.h&gt;

24493 size\_t mbsrtowcs(wchar\_t \*restrict *dst*, const char \*\*restrict *src*,  
24494 size\_t *len*, mbstate\_t \*restrict *ps*);24495 **DESCRIPTION**24496 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
24497 conflict between the requirements described here and the ISO C standard is unintentional. This  
24498 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.24499 The *mbsrtowcs()* function shall convert a sequence of characters, beginning in the conversion  
24500 state described by the object pointed to by *ps*, from the array indirectly pointed to by *src* into a  
24501 sequence of corresponding wide characters. If *dst* is not a null pointer, the converted characters  
24502 are stored into the array pointed to by *dst*. Conversion continues up to and including a  
24503 terminating null character, which is also stored. Conversion stops early in either of the following  
24504 cases:

- 24505
- A sequence of bytes is encountered that does not form a valid character.
  - *len* codes have been stored into the array pointed to by *dst* (and *dst* is not a null pointer).
- 24506

24507 Each conversion takes place as if by a call to the *mbrtowc()* function.24508 If *dst* is not a null pointer, the pointer object pointed to by *src* is assigned either a null pointer (if  
24509 conversion stopped due to reaching a terminating null character) or the address just past the last  
24510 character converted (if any). If conversion stopped due to reaching a terminating null character,  
24511 and if *dst* is not a null pointer, the resulting state described is the initial conversion state.24512 If *ps* is a null pointer, the *mbsrtowcs()* function uses its own internal **mbstate\_t** object, which is  
24513 initialized at program start-up to the initial conversion state. Otherwise, the **mbstate\_t** object  
24514 pointed to by *ps* is used to completely describe the current conversion state of the associated  
24515 character sequence. The implementation behaves as if no function defined in this volume of  
24516 IEEE Std. 1003.1-200x calls *mbsrtowcs()*.24517 **XSI** The behavior of this function shall be affected by the *LC\_CTYPE* category of the current locale.24518 **RETURN VALUE**24519 If the input conversion encounters a sequence of bytes that do not form a valid character, an  
24520 encoding error occurs. In this case, the *mbsrtowcs()* function stores the value of the macro  
24521 [EILSEQ] in *errno* and shall return (**size\_t**)-1; the conversion state is undefined. Otherwise, it  
24522 shall return the number of characters successfully converted, not including the terminating null  
24523 (if any).24524 **ERRORS**24525 The *mbsrtowcs()* function may fail if:

- 24526 [EINVAL]
- ps*
- points to an object that contains an invalid conversion state.
- 
- 24527 [EILSEQ] Invalid character sequence is detected.

24528 **EXAMPLES**

24529 None.

24530 **APPLICATION USAGE**

24531 None.

24532 **RATIONALE**

24533 None.

24534 **FUTURE DIRECTIONS**

24535 None.

24536 **SEE ALSO**24537 *mbsinit()*, *mbrtowc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`24538 **CHANGE HISTORY**24539 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
24540 (E).24541 **Issue 6**24542 The *mbsrtowcs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

24543 **NAME**

24544 mbstowcs — convert a character string to a wide-character string

24545 **SYNOPSIS**

24546 #include <stdlib.h>

24547 size\_t mbstowcs(wchar\_t \*restrict pwcs, const char \*restrict s,  
24548 size\_t n);

24549 **DESCRIPTION**

24550 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
24551 conflict between the requirements described here and the ISO C standard is unintentional. This  
24552 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24553 The *mbstowcs()* function shall convert a sequence of characters that begins in the initial shift  
24554 state from the array pointed to by *s* into a sequence of corresponding wide-character codes and  
24555 stores not more than *n* wide-character codes into the array pointed to by *pwcs*. No characters  
24556 that follow a null byte (which is converted into a wide-character code with value 0) shall be  
24557 examined or converted. Each character is converted as if by a call to *mbtowc()*, except that the  
24558 shift state of *mbtowc()* is not affected.

24559 No more than *n* elements shall be modified in the array pointed to by *pwcs*. If copying takes  
24560 place between objects that overlap, the behavior is undefined.

24561 XSI The behavior of this function shall be affected by the *LC\_CTYPE* category of the current locale. If  
24562 *pwcs* is a null pointer, *mbstowcs()* shall return the length required to convert the entire array  
24563 regardless of the value of *n*, but no values are stored.

24564 **RETURN VALUE**

24565 CX If an invalid character is encountered, *mbstowcs()* shall return **(size\_t)-1** and may set *errno* to  
24566 indicate the error. Otherwise, *mbstowcs()* shall return the number of the array elements modified  
24567 (or required if *pwcs* is null), not including a terminating 0 code, if any. The array shall not be  
24568 zero-terminated if the value returned is *n*.

24569 **ERRORS**

24570 The *mbstowcs()* function may fail if:

24571 XSI [EILSEQ] Invalid byte sequence is detected.

24572 **EXAMPLES**

24573 None.

24574 **APPLICATION USAGE**

24575 None.

24576 **RATIONALE**

24577 None.

24578 **FUTURE DIRECTIONS**

24579 None.

24580 **SEE ALSO**

24581 *mblen()*, *mbtowc()*, *wctomb()*, *wcstombs()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
24582 <stdlib.h>

24583 **CHANGE HISTORY**

24584 First released in Issue 4. Aligned with the ISO C standard.

24585 **Issue 6**

24586

The *mbstowcs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

24587 **NAME**

24588           mbtowc — convert a character to a wide-character code

24589 **SYNOPSIS**

24590           #include &lt;stdlib.h&gt;

24591           int mbtowc(wchar\_t \*restrict *pwc*, const char \*restrict *s*, size\_t *n*);24592 **DESCRIPTION**

24593 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 24594 conflict between the requirements described here and the ISO C standard is unintentional. This  
 24595 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24596       If *s* is not a null pointer, *mbtowc()* shall determine the number of the bytes that constitute the  
 24597 character pointed to by *s*. It then determines the wide-character code for the value of type  
 24598 **wchar\_t** that corresponds to that character. (The value of the wide-character code corresponding  
 24599 to the null byte is 0.) If the character is valid and *pwc* is not a null pointer, *mbtowc()* stores the  
 24600 wide-character code in the object pointed to by *pwc*.

24601       The behavior of this function is affected by the *LC\_CTYPE* category of the current locale. For a  
 24602 state-dependent encoding, this function is placed into its initial state by a call for which its  
 24603 character pointer argument, *s*, is a null pointer. Subsequent calls with *s* as other than a null  
 24604 pointer cause the internal state of the function to be altered as necessary. A call with *s* as a null  
 24605 pointer causes this function to return a non-zero value if encodings have state dependency, and  
 24606 0 otherwise. If the implementation employs special bytes to change the shift state, these bytes do  
 24607 not produce separate wide-character codes, but are grouped with an adjacent character.  
 24608 Changing the *LC\_CTYPE* category causes the shift state of this function to be indeterminate. At  
 24609 most *n* bytes of the array pointed to by *s* shall be examined.

24610       The implementation shall behave as if no function defined in this volume of  
 24611 IEEE Std. 1003.1-200x calls *mbtowc()*.

24612 **RETURN VALUE**

24613       If *s* is a null pointer, *mbtowc()* shall return a non-zero or 0 value, if character encodings,  
 24614 respectively, do or do not have state-dependent encodings. If *s* is not a null pointer, *mbtowc()*  
 24615 shall either return 0 (if *s* points to the null byte), or return the number of bytes that constitute the  
 24616 **CX** converted character (if the next *n* or fewer bytes form a valid character), or return -1 and may  
 24617 set *errno* to indicate the error (if they do not form a valid character).

24618       In no case shall the value returned be greater than *n* or the value of the {MB\_CUR\_MAX} macro.

24619 **ERRORS**

24620       The *mbtowc()* function may fail if:

24621 **XSI**       [EILSEQ]       Invalid character sequence is detected.

24622 **EXAMPLES**

24623       None.

24624 **APPLICATION USAGE**

24625       None.

24626 **RATIONALE**

24627       None.

24628 **FUTURE DIRECTIONS**

24629       None.

24630 **SEE ALSO**

24631 *mblen()*, *mbstowcs()*, *wctomb()*, *wcstombs()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
24632 **<stdlib.h>**

24633 **CHANGE HISTORY**

24634 First released in Issue 4. Aligned with the ISO C standard.

24635 **Issue 6**

24636 The *mbtowc()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

24637 **NAME**

24638 memccpy — copy bytes in memory

24639 **SYNOPSIS**

24640 XSI #include &lt;string.h&gt;

24641 void \*memccpy(void \*restrict *s1*, const void \*restrict *s2*,  
24642 int *c*, size\_t *n*);

24643

24644 **DESCRIPTION**

24645 The *memccpy()* function shall copy bytes from memory area *s2* into *s1*, stopping after the first  
24646 occurrence of byte *c* (converted to an **unsigned char**) is copied, or after *n* bytes are copied,  
24647 whichever comes first. If copying takes place between objects that overlap, the behavior is  
24648 undefined.

24649 **RETURN VALUE**

24650 The *memccpy()* function shall return a pointer to the byte after the copy of *c* in *s1*, or a null  
24651 pointer if *c* was not found in the first *n* bytes of *s2*.

24652 **ERRORS**

24653 No errors are defined.

24654 **EXAMPLES**

24655 None.

24656 **APPLICATION USAGE**24657 The *memccpy()* function does not check for the overflow of the receiving memory area.24658 **RATIONALE**

24659 None.

24660 **FUTURE DIRECTIONS**

24661 None.

24662 **SEE ALSO**

24663 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;string.h&gt;

24664 **CHANGE HISTORY**

24665 First released in Issue 1. Derived from Issue 1 of the SVID.

24666 **Issue 4**24667 The type of argument *s2* is changed from **void\*** to **const void\***.

24668 Reference to use of the <**memory.h**> header is removed from the APPLICATION USAGE  
24669 section.

24670 The FUTURE DIRECTIONS section is removed.

24671 **Issue 6**

24672 The **restrict** keyword is added to the *memccpy()* prototype for alignment with the  
24673 ISO/IEC 9899:1999 standard.

24674 **NAME**

24675 memchr — find byte in memory

24676 **SYNOPSIS**

24677 #include <string.h>

24678 void \*memchr(const void \*s, int c, size\_t n);

24679 **DESCRIPTION**

24680 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
24681 conflict between the requirements described here and the ISO C standard is unintentional. This  
24682 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24683 The *memchr()* function shall locate the first occurrence of *c* (converted to an **unsigned char**) in  
24684 the initial *n* bytes (each interpreted as **unsigned char**) of the object pointed to by *s*.

24685 **RETURN VALUE**

24686 The *memchr()* function shall return a pointer to the located byte, or a null pointer if the byte does  
24687 not occur in the object.

24688 **ERRORS**

24689 No errors are defined.

24690 **EXAMPLES**

24691 None.

24692 **APPLICATION USAGE**

24693 None.

24694 **RATIONALE**

24695 None.

24696 **FUTURE DIRECTIONS**

24697 None.

24698 **SEE ALSO**

24699 The Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>

24700 **CHANGE HISTORY**

24701 First released in Issue 1. Derived from Issue 1 of the SVID.

24702 **Issue 4**

24703 The APPLICATION USAGE section is removed.

24704 The following changes are incorporated for alignment with the ISO C standard:

- 24705 • The function is no longer marked as an extension.
- 24706 • The type of argument *s* is changed from **void\*** to **const void\***.

24707 **NAME**24708        **memcmp** — compare bytes in memory24709 **SYNOPSIS**

24710        #include &lt;string.h&gt;

24711        int memcmp(const void \*s1, const void \*s2, size\_t n);

24712 **DESCRIPTION**

24713 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
24714        conflict between the requirements described here and the ISO C standard is unintentional. This  
24715        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24716        The *memcmp()* function shall compare the first *n* bytes (each interpreted as **unsigned char**) of the  
24717        object pointed to by *s1* to the first *n* bytes of the object pointed to by *s2*.

24718        The sign of a non-zero return value shall be determined by the sign of the difference between the  
24719        values of the first pair of bytes (both interpreted as type **unsigned char**) that differ in the objects  
24720        being compared.

24721 **RETURN VALUE**

24722        The *memcmp()* function shall return an integer greater than, equal to, or less than 0, if the object  
24723        pointed to by *s1* is greater than, equal to, or less than the object pointed to by *s2*, respectively.

24724 **ERRORS**

24725        No errors are defined.

24726 **EXAMPLES**

24727        None.

24728 **APPLICATION USAGE**

24729        None.

24730 **RATIONALE**

24731        None.

24732 **FUTURE DIRECTIONS**

24733        None.

24734 **SEE ALSO**24735        The Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>24736 **CHANGE HISTORY**

24737        First released in Issue 1. Derived from Issue 1 of the SVID.

24738 **Issue 4**

24739        The RETURN VALUE section is clarified.

24740        The APPLICATION USAGE section is removed.

24741        The following changes are incorporated for alignment with the ISO C standard:

- 24742        • The function is no longer marked as an extension.
- 24743        • The type of arguments *s1* and *s2* are changed from **void\*** to **const void\***.

24744 **NAME**

24745           memcpy — copy bytes in memory

24746 **SYNOPSIS**

24747           #include <string.h>

24748           void \*memcpy(void \*restrict *s1*, const void \*restrict *s2*, size\_t *n*);

24749 **DESCRIPTION**

24750 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
24751           conflict between the requirements described here and the ISO C standard is unintentional. This  
24752           volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24753           The *memcpy()* function shall copy *n* bytes from the object pointed to by *s2* into the object pointed  
24754           to by *s1*. If copying takes place between objects that overlap, the behavior is undefined.

24755 **RETURN VALUE**

24756           The *memcpy()* function shall return *s1*; no return value is reserved to indicate an error.

24757 **ERRORS**

24758           No errors are defined.

24759 **EXAMPLES**

24760           None.

24761 **APPLICATION USAGE**

24762           The *memcpy()* function does not check for the overflowing of the receiving memory area.

24763 **RATIONALE**

24764           None.

24765 **FUTURE DIRECTIONS**

24766           None.

24767 **SEE ALSO**

24768           The Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>

24769 **CHANGE HISTORY**

24770           First released in Issue 1. Derived from Issue 1 of the SVID.

24771 **Issue 4**

24772           Reference to use of the <**memory.h**> header is removed from the APPLICATION USAGE  
24773           section, and a note about overflow checking has been added.

24774           The FUTURE DIRECTIONS section is removed.

24775           The following changes are incorporated for alignment with the ISO C standard:

- 24776           • The function is no longer marked as an extension.
- 24777           • The type of argument *s2* is changed from **void\*** to **const void\***.

24778 **Issue 6**

24779           The *memcpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

24780 **NAME**

24781 memmove — copy bytes in memory with overlapping areas

24782 **SYNOPSIS**

24783 #include &lt;string.h&gt;

24784 void \*memmove(void \*s1, const void \*s2, size\_t n);

24785 **DESCRIPTION**

24786 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
24787 conflict between the requirements described here and the ISO C standard is unintentional. This  
24788 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24789 The *memmove()* function shall copy *n* bytes from the object pointed to by *s2* into the object  
24790 pointed to by *s1*. Copying takes place as if the *n* bytes from the object pointed to by *s2* are first  
24791 copied into a temporary array of *n* bytes that does not overlap the objects pointed to by *s1* and  
24792 *s2*, and then the *n* bytes from the temporary array are copied into the object pointed to by *s1*.

24793 **RETURN VALUE**24794 The *memmove()* function shall return *s1*; no return value is reserved to indicate an error.24795 **ERRORS**

24796 No errors are defined.

24797 **EXAMPLES**

24798 None.

24799 **APPLICATION USAGE**

24800 None.

24801 **RATIONALE**

24802 None.

24803 **FUTURE DIRECTIONS**

24804 None.

24805 **SEE ALSO**

24806 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;string.h&gt;

24807 **CHANGE HISTORY**

24808 First released in Issue 4. Derived from the ANSI C standard.

24809 **NAME**

24810           memset — set bytes in memory

24811 **SYNOPSIS**

24812           #include <string.h>

24813           void \*memset(void \*s, int c, size\_t n);

24814 **DESCRIPTION**

24815 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
24816       conflict between the requirements described here and the ISO C standard is unintentional. This  
24817       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

24818       The *memset()* function shall copy *c* (converted to an **unsigned char**) into each of the first *n* bytes  
24819       of the object pointed to by *s*.

24820 **RETURN VALUE**

24821       The *memset()* function shall return *s*; no return value is reserved to indicate an error.

24822 **ERRORS**

24823       No errors are defined.

24824 **EXAMPLES**

24825       None.

24826 **APPLICATION USAGE**

24827       None.

24828 **RATIONALE**

24829       None.

24830 **FUTURE DIRECTIONS**

24831       None.

24832 **SEE ALSO**

24833       The Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>

24834 **CHANGE HISTORY**

24835       First released in Issue 1. Derived from Issue 1 of the SVID.

24836 **Issue 4**

24837       The APPLICATION USAGE section is removed.

24838       The following change is incorporated for alignment with the ISO C standard:

- 24839
  - The function is no longer marked as an extension.

## 24840 NAME

24841 mkdir — make a directory

## 24842 SYNOPSIS

24843 #include &lt;sys/stat.h&gt;

24844 int mkdir(const char \*path, mode\_t mode);

## 24845 DESCRIPTION

24846 The *mkdir()* function creates a new directory with name *path*. The file permission bits of the new  
 24847 directory are initialized from *mode*. These file permission bits of the *mode* argument are modified  
 24848 by the process' file creation mask.

24849 When bits in *mode* other than the file permission bits are set, the meaning of these additional bits  
 24850 is implementation-defined.

24851 The directory's user ID is set to the process' effective user ID. The directory's group ID shall be  
 24852 set to the group ID of the parent directory or to the effective group ID of the process.

24853 The newly created directory shall be an empty directory.

24854 If *path* names a symbolic link, *mkdir()* shall fail and set *errno* to [EEXIST].

24855 Upon successful completion, *mkdir()* shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime*  
 24856 fields of the directory. Also, the *st\_ctime* and *st\_mtime* fields of the directory that contains the  
 24857 new entry shall be marked for update.

## 24858 RETURN VALUE

24859 Upon successful completion, *mkdir()* shall return 0. Otherwise, -1 shall be returned, no directory  
 24860 shall be created, and *errno* shall be set to indicate the error.

## 24861 ERRORS

24862 The *mkdir()* function shall fail if:

24863 [EACCES] Search permission is denied on a component of the path prefix, or write  
 24864 permission is denied on the parent directory of the directory to be created.

24865 [EEXIST] The named file exists.

24866 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 24867 argument.

24868 [EMLINK] The link count of the parent directory would exceed {LINK\_MAX}.

24869 [ENAMETOOLONG]

24870 The length of the *path* argument exceeds {PATH\_MAX} or a path name  
 24871 component is longer than {NAME\_MAX}.

24872 [ENOENT] A component of the path prefix specified by *path* does not name an existing  
 24873 directory or *path* is an empty string.

24874 [ENOSPC] The file system does not contain enough space to hold the contents of the new  
 24875 directory or to extend the parent directory of the new directory.

24876 [ENOTDIR] A component of the path prefix is not a directory.

24877 [EROFS] The parent directory resides on a read-only file system.

24878 The *mkdir()* function may fail if:

24879 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 24880 resolution of the *path* argument.

24881 [ENAMETOOLONG]

24882 As a result of encountering a symbolic link in resolution of the *path* argument,  
24883 the length of the substituted path name string exceeded {PATH\_MAX}.

#### 24884 EXAMPLES

##### 24885 Creating a Directory

24886 The following example shows how to create a directory named `/home/cnd/mod1`, with  
24887 read/write/search permissions for owner and group, and with read/search permissions for  
24888 others.

```
24889 #include <sys/types.h>
```

```
24890 #include <sys/stat.h>
```

```
24891 int status;
```

```
24892 ...
```

```
24893 status = mkdir("/home/cnd/mod1", S_IRWXU | S_IRWXG | S_IROTH | S_IXOTH);
```

#### 24894 APPLICATION USAGE

24895 None.

#### 24896 RATIONALE

24897 The `mkdir()` function originated in 4.2 BSD and was added to System V in Release 3.0.

24898 4.3 BSD detects [ENAMETOOLONG].

24899 See `getgroups()` about the group of a newly created directory.

#### 24900 FUTURE DIRECTIONS

24901 None.

#### 24902 SEE ALSO

24903 `umask()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<sys/stat.h>`, `<sys/types.h>`

#### 24904 CHANGE HISTORY

24905 First released in Issue 3.

24906 Entry included for alignment with the POSIX.1-1988 standard.

#### 24907 Issue 4

24908 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on  
24909 XSI-conformant systems.

24910 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 24911 • The type of argument *path* is changed from `char*` to `const char*`.

24912 The following changes are incorporated for alignment with the FIPS requirements:

- 24913 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
24914 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
24915 an extension.

#### 24916 Issue 4, Version 2

24917 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 24918 • It states that [ELOOP] is returned if too many symbolic links are encountered during path  
24919 name resolution.
- 24920 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an  
24921 intermediate result of path name resolution of a symbolic link.

24922 **Issue 6**

24923 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

24924 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 24925 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.
- 24926 This is since behavior may vary from one file system to another.

24927 The following new requirements on POSIX implementations derive from alignment with the  
24928 Single UNIX Specification:

- 24929 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
24930 required for conforming implementations of previous POSIX specifications, it was not  
24931 required for UNIX applications.
- 24932 • The [ELOOP] mandatory error condition is added.
- 24933 • A second [ENAMETOOLONG] is added as an optional error condition.

24934 The following changes were made to align with the IEEE P1003.1a draft standard:

- 24935 • The [ELOOP] optional error condition is added.

24936 **NAME**

24937 mkfifo — make a FIFO special file

24938 **SYNOPSIS**

24939 #include &lt;sys/stat.h&gt;

24940 int mkfifo(const char \*path, mode\_t mode);

24941 **DESCRIPTION**

24942 The *mkfifo()* function shall create a new FIFO special file named by the path name pointed to by  
 24943 *path*. The file permission bits of the new FIFO are initialized from *mode*. The file permission bits  
 24944 of the *mode* argument are modified by the process' file creation mask.

24945 When bits in *mode* other than the file permission bits are set, the effect is implementation-  
 24946 defined.

24947 If *path* names a symbolic link, *mkfifo()* shall fail and set *errno* to [EEXIST].

24948 The FIFO's user ID shall be set to the process' effective user ID. The FIFO's group ID shall be set  
 24949 to the group ID of the parent directory or to the effective group ID of the process.

24950 Upon successful completion, *mkfifo()* shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime*  
 24951 fields of the file. Also, the *st\_ctime* and *st\_mtime* fields of the directory that contains the new  
 24952 entry shall be marked for update.

24953 **RETURN VALUE**

24954 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned, no FIFO shall  
 24955 be created, and *errno* shall be set to indicate the error.

24956 **ERRORS**24957 The *mkfifo()* function shall fail if:

24958 [EACCES] A component of the path prefix denies search permission, or write permission  
 24959 is denied on the parent directory of the FIFO to be created.

24960 [EEXIST] The named file already exists.

24961 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 24962 argument.

24963 [ENAMETOOLONG]

24964 The length of the *path* argument exceeds {PATH\_MAX} or a path name  
 24965 component is longer than {NAME\_MAX}.

24966 [ENOENT] A component of the path prefix specified by *path* does not name an existing  
 24967 directory or *path* is an empty string.

24968 [ENOSPC] The directory that would contain the new file cannot be extended or the file  
 24969 system is out of file-allocation resources.

24970 [ENOTDIR] A component of the path prefix is not a directory.

24971 [EROFS] The named file resides on a read-only file system.

24972 The *mkfifo()* function may fail if:

24973 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 24974 resolution of the *path* argument.

24975 [ENAMETOOLONG]

24976 As a result of encountering a symbolic link in resolution of the *path* argument,  
 24977 the length of the substituted path name string exceeded {PATH\_MAX}.

24978 **EXAMPLES**24979 **Creating a FIFO File**

24980 The following example shows how to create a FIFO file named `/home/cnd/mod_done`, with  
 24981 read/write permissions for owner, and with read permissions for group and others.

```
24982 #include <sys/types.h>
24983 #include <sys/stat.h>
24984 int status;
24985 ...
24986 status = mkfifo("/home/cnd/mod_done", S_IWUSR | S_IRUSR |
24987 S_IRGRP | S_IROTH);
```

24988 **APPLICATION USAGE**

24989 None.

24990 **RATIONALE**

24991 The syntax of this function is intended to maintain compatibility with historical  
 24992 implementations of `mknod()`. The latter function was included in the 1984 `/usr/group` standard  
 24993 but only for use in creating FIFO special files. The `mknod()` function was originally excluded  
 24994 from the POSIX.1-1988 standard as implementation-defined and replaced by `mkdir()` and  
 24995 `mkfifo()`. The `mknod()` function is now included for alignment with the Single UNIX  
 24996 Specification.

24997 See `getgroups()` about the group of a newly created FIFO.

24998 **FUTURE DIRECTIONS**

24999 None.

25000 **SEE ALSO**

25001 `umask()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<sys/stat.h>`, `<sys/types.h>`

25002 **CHANGE HISTORY**

25003 First released in Issue 3.

25004 Entry included for alignment with the POSIX.1-1988 standard.

25005 **Issue 4**

25006 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on  
 25007 XSI-conformant systems.

25008 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 25009 • The type of argument `path` is changed from `char*` to `const char*`.
- 25010 • The description of [EACCES] is updated to indicate that this error is also returned if write  
 25011 permission is denied to the parent directory.

25012 The following changes are incorporated for alignment with the FIPS requirements:

- 25013 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
 25014 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
 25015 an extension.

25016 **Issue 4, Version 2**

25017 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 25018 • It states that [ELOOP] is returned if too many symbolic links are encountered during path  
 25019 name resolution.

- 25020           • A second [ENAMETOOLONG] condition is defined that may report excessive length of an  
25021           intermediate result of path name resolution of a symbolic link.
- 25022 **Issue 6**
- 25023           In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.
- 25024           The following changes are made for alignment with the ISO POSIX-1: 1996 standard:
- 25025           • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.  
25026           This is since behavior may vary from one file system to another.
- 25027           The following new requirements on POSIX implementations derive from alignment with the  
25028           Single UNIX Specification:
- 25029           • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
25030           required for conforming implementations of previous POSIX specifications, it was not  
25031           required for UNIX applications.
- 25032           • The [ELOOP] mandatory error condition is added.
- 25033           • A second [ENAMETOOLONG] is added as an optional error condition.
- 25034           The following changes were made to align with the IEEE P1003.1a draft standard:
- 25035           • The [ELOOP] optional error condition is added.

25036 **NAME**

25037 `mknod` — make a directory, a special or regular file

25038 **SYNOPSIS**

25039 XSI `#include <sys/stat.h>`

25040 `int mknod(const char *path, mode_t mode, dev_t dev);`

25041

25042 **DESCRIPTION**

25043 The `mknod()` function shall create a new file named by the path name to which the argument  
25044 `path` points.

25045 The file type for `path` is OR'ed into the `mode` argument, and the application shall select one of the  
25046 following symbolic constants:

25047

25048

25049

25050

25051

25052

25053

| Name    | Description                      |
|---------|----------------------------------|
| S_IFIFO | FIFO-special                     |
| S_IFCHR | Character-special (non-portable) |
| S_IFDIR | Directory (non-portable)         |
| S_IFBLK | Block-special (non-portable)     |
| S_IFREG | Regular (non-portable)           |

25054

25055

The only portable use of `mknod()` is to create a FIFO-special file. If `mode` is not S\_IFIFO or `dev` is not 0, the behavior of `mknod()` is unspecified.

25056

25057

The permissions for the new file are OR'ed into the `mode` argument, and may be selected from any combination of the following symbolic constants:

25058

25059

25060

25061

25062

25063

25064

25065

25066

25067

25068

25069

25070

25071

25072

25073

25074

| Name    | Description                                 |
|---------|---------------------------------------------|
| S_ISUID | Set user ID on execution.                   |
| S_ISGID | Set group ID on execution.                  |
| S_IRWXU | Read, write, or execute (search) by owner.  |
| S_IRUSR | Read by owner.                              |
| S_IWUSR | Write by owner.                             |
| S_IXUSR | Execute (search) by owner.                  |
| S_IRWXG | Read, write, or execute (search) by group.  |
| S_IRGRP | Read by group.                              |
| S_IWGRP | Write by group.                             |
| S_IXGRP | Execute (search) by group.                  |
| S_IRWXO | Read, write, or execute (search) by others. |
| S_IROTH | Read by others.                             |
| S_IWOTH | Write by others.                            |
| S_IXOTH | Execute (search) by others.                 |
| S_ISVTX | On directories, restricted deletion flag.   |

25075

25076

25077

The user ID of the file is initialized to the effective user ID of the process. The group ID of the file is initialized to either the effective group ID of the process or the group ID of the parent directory.

25078

25079

25080

The owner, group, and other permission bits of `mode` are modified by the file mode creation mask of the process. The `mknod()` function clears each bit whose corresponding bit in the file mode creation mask of the process is set.

- 25081 If *path* names a symbolic link, *mknod()* shall fail and set *errno* to [EEXIST].
- 25082 Upon successful completion, *mknod()* shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime*  
25083 fields of the file. Also, the *st\_ctime* and *st\_mtime* fields of the directory that contains the new  
25084 entry shall be marked for update.
- 25085 Only a process with appropriate privileges may invoke *mknod()* for file types other than FIFO-  
25086 special.
- 25087 **RETURN VALUE**
- 25088 Upon successful completion, *mknod()* shall return 0. Otherwise, it shall return -1, the new file  
25089 shall not be created, and *errno* shall be set to indicate the error.
- 25090 **ERRORS**
- 25091 The *mknod()* function shall fail if:
- 25092 [EACCES] A component of the path prefix denies search permission, or write permission  
25093 is denied on the parent directory.
- 25094 [EEXIST] The named file exists.
- 25095 [EINVAL] An invalid argument exists.
- 25096 [EIO] An I/O error occurred while accessing the file system.
- 25097 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
25098 argument.
- 25099 [ENAMETOOLONG]
- 25100 The length of a path name exceeds {PATH\_MAX} or a path name component  
25101 is longer than {NAME\_MAX}.
- 25102 [ENOENT] A component of the path prefix specified by *path* does not name an existing  
25103 directory or *path* is an empty string.
- 25104 [ENOSPC] The directory that would contain the new file cannot be extended or the file  
25105 system is out of file allocation resources.
- 25106 [ENOTDIR] A component of the path prefix is not a directory.
- 25107 [EPERM] The invoking process does not have appropriate privileges and the file type is  
25108 not FIFO-special.
- 25109 [EROFS] The directory in which the file is to be created is located on a read-only file  
25110 system.
- 25111 The *mknod()* function may fail if:
- 25112 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
25113 resolution of the *path* argument.
- 25114 [ENAMETOOLONG]
- 25115 Path name resolution of a symbolic link produced an intermediate result  
25116 whose length exceeds {PATH\_MAX}.

25117 **EXAMPLES**25118 **Creating a FIFO Special File**

25119 The following example shows how to create a FIFO special file named `/home/cnd/mod_done`,  
25120 with read/write permissions for owner, and with read permissions for group and others.

```
25121 #include <sys/types.h>
25122 #include <sys/stat.h>
25123 dev_t dev;
25124 int status;
25125 ...
25126 status = mknod("/home/cnd/mod_done", S_IFIFO | S_IWUSR |
25127 S_IRUSR | S_IRGRP | S_IROTH, dev);
```

25128 **APPLICATION USAGE**

25129 *mkfifo()* is preferred over this function for making FIFO special files.

25130 **RATIONALE**

25131 None.

25132 **FUTURE DIRECTIONS**

25133 None.

25134 **SEE ALSO**

25135 *chmod()*, *creat()*, *exec*, *mkdir()*, *mkfifo()*, *open()*, *stat()*, *umask()*, the Base Definitions volume of  
25136 IEEE Std. 1003.1-200x, `<sys/stat.h>`

25137 **CHANGE HISTORY**

25138 First released in Issue 4, Version 2.

25139 **Issue 5**

25140 Moved from X/OPEN UNIX extension to BASE.

25141 **Issue 6**

25142 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

25143 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
25144 [ELOOP] error condition is added.

25145 **NAME**

25146 mkstemp — make a unique file name

25147 **SYNOPSIS**25148 XSI `#include <stdlib.h>`25149 `int mkstemp(char *template);`

25150

25151 **DESCRIPTION**

25152 The *mkstemp()* function shall replace the contents of the string pointed to by *template* by a unique  
25153 file name, and return a file descriptor for the file open for reading and writing. The function thus  
25154 prevents any possible race condition between testing whether the file exists and opening it for  
25155 use. The string in *template* should look like a file name with six trailing *Xs*; *mkstemp()* replaces  
25156 each *X* with a character from the portable file name character set. The characters are chosen such  
25157 that the resulting name does not duplicate the name of an existing file at the time of a call to  
25158 *mkstemp()*.

25159 **RETURN VALUE**

25160 Upon successful completion, *mkstemp()* shall return an open file descriptor. Otherwise,  $-1$  shall  
25161 be returned if no suitable file could be created.

25162 **ERRORS**

25163 No errors are defined.

25164 **EXAMPLES**25165 **Generating a File Name**

25166 The following example creates a file with a 10-character name beginning with the characters  
25167 "file" and opens the file for reading and writing. The value returned as the value of *fd* is a file  
25168 descriptor that identifies the file.

```
25169 #include <stdlib.h>
25170 ...
25171 char *template = "/tmp/fileXXXXXX";
25172 int fd;
25173 fd = mkstemp(template);
```

25174 **APPLICATION USAGE**

25175 It is possible to run out of letters.

25176 The *mkstemp()* function need not check to determine whether the file name part of *template*  
25177 exceeds the maximum allowable file name length.

25178 **RATIONALE**

25179 None.

25180 **FUTURE DIRECTIONS**

25181 None.

25182 **SEE ALSO**

25183 *getpid()*, *open()*, *tmpfile()*, *tmpnam()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
25184 `<stdlib.h>`

25185 **CHANGE HISTORY**

25186           First released in Issue 4, Version 2.

25187 **Issue 5**

25188           Moved from X/OPEN UNIX extension to BASE.

25189 **NAME**25190 mktemp — make a unique file name (**LEGACY**)25191 **SYNOPSIS**25192 XSI `#include <stdlib.h>`25193 `char *mktemp(char *template);`

25194

25195 **DESCRIPTION**

25196 The *mktemp()* function shall replace the contents of the string pointed to by *template* by a unique  
 25197 file name and return *template*. The application shall initialize *template* to be a file name with six  
 25198 trailing *Xs*; *mktemp()* shall replace each *X* with a single byte character from the portable file  
 25199 name character set.

25200 **RETURN VALUE**

25201 The *mktemp()* function shall return the pointer *template*. If a unique name cannot be created,  
 25202 *template* shall point to a null string.

25203 **ERRORS**

25204 No errors are defined.

25205 **EXAMPLES**25206 **Generating a File Name**

25207 The following example replaces the contents of the "template" string with a 10-character file  
 25208 name beginning with the characters "file" and returns a pointer to the "template" string  
 25209 that contains the new file name.

25210 `#include <stdlib.h>`25211 `...`25212 `char *template = "/tmp/fileXXXXXX";`25213 `char *ptr;`25214 `ptr = mktemp(template);`25215 **APPLICATION USAGE**

25216 Between the time a path name is created and the file opened, it is possible for some other process  
 25217 to create a file with the same name. The *mkstemp()* function avoids this problem and is preferred  
 25218 over this function.

25219 **RATIONALE**

25220 None.

25221 **FUTURE DIRECTIONS**

25222 This function may be withdrawn in a future version.

25223 **SEE ALSO**25224 *mkstemp()*, *tmpfile()*, *tmpnam()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>` |25225 **CHANGE HISTORY**

25226 First released in Issue 4, Version 2.

25227 **Issue 5**

25228 Moved from X/OPEN UNIX extension to BASE.

25229 **Issue 6**

25230 This function is marked LEGACY.

25231 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

25232 **NAME**

25233 mktime — convert broken-down time into time since the Epoch

25234 **SYNOPSIS**

25235 #include &lt;time.h&gt;

25236 time\_t mktime(struct tm \*timeptr);

25237 **DESCRIPTION**

25238 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 25239 conflict between the requirements described here and the ISO C standard is unintentional. This  
 25240 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

25241 The *mktime()* function shall convert the broken-down time, expressed as local time, in the  
 25242 structure pointed to by *timeptr*, into a time since the Epoch value with the same encoding as that  
 25243 of the values returned by *time()*. The original values of the *tm\_wday* and *tm\_yday* components of  
 25244 the structure are ignored, and the original values of the other components are not restricted to  
 25245 the ranges described in <time.h>.

25246 **CX** A positive or 0 value for *tm\_isdst* shall cause *mktime()* to presume initially that Daylight Savings  
 25247 Time, respectively, is or is not in effect for the specified time. A negative value for *tm\_isdst* shall  
 25248 cause *mktime()* to attempt to determine whether Daylight Saving Time is in effect for the  
 25249 specified time.

25250 Local timezone information shall be set as though *mktime()* called *tzset()*.

25251 Upon successful completion, the values of the *tm\_wday* and *tm\_yday* components of the structure  
 25252 shall be set appropriately, and the other components are set to represent the specified time since  
 25253 the Epoch, but with their values forced to the ranges indicated in the <time.h> entry; the final  
 25254 value of *tm\_mday* shall not be set until *tm\_mon* and *tm\_year* are determined.

25255 **RETURN VALUE**

25256 The *mktime()* function shall return the specified time since the Epoch encoded as a value of type  
 25257 **time\_t**. If the time since the Epoch cannot be represented, the function shall return the value  
 25258 **(time\_t)-1**.

25259 **ERRORS**

25260 No errors are defined.

25261 **EXAMPLES**

25262 What day of the week is July 4, 2001?

25263 #include &lt;stdio.h&gt;

25264 #include &lt;time.h&gt;

25265 struct tm time\_str;

25266 char daybuf[20];

25267 int main(void)

25268 {

25269 time\_str.tm\_year = 2001 - 1900;

25270 time\_str.tm\_mon = 7 - 1;

25271 time\_str.tm\_mday = 4;

25272 time\_str.tm\_hour = 0;

25273 time\_str.tm\_min = 0;

25274 time\_str.tm\_sec = 1;

25275 time\_str.tm\_isdst = -1;

25276 if (mktime(&amp;time\_str) == -1)

```
25277 (void)puts("-unknown-");
25278 else {
25279 (void)strftime(daybuf, sizeof(daybuf), "%A", &time_str);
25280 (void)puts(daybuf);
25281 }
25282 return 0;
25283 }
```

**25284 APPLICATION USAGE**

25285 None.

**25286 RATIONALE**

25287 None.

**25288 FUTURE DIRECTIONS**

25289 None.

**25290 SEE ALSO**

25291 *asctime()*, *clock()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *strftime()*, *strptime()*, *time()*, *utime()*,  
25292 the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

**25293 CHANGE HISTORY**

25294 First released in Issue 3.

25295 Entry included for alignment with the POSIX.1-1988 standard and the ANSI C standard.

**25296 Issue 4**

25297 In the DESCRIPTION, a paragraph is added indicating the possible settings of *tm\_isdst*, and  
25298 reference to setting of *tm\_sec* for leap seconds or double leap seconds is removed (although this  
25299 functionality is still supported).

25300 In the EXAMPLES section, the sample code is updated to use ISO C standard syntax.

**25301 Issue 6**

25302 Extensions beyond the ISO C standard are now marked.

25303 **NAME**25304 mlock, munlock — lock or unlock a range of process address space (**REALTIME**)25305 **SYNOPSIS**

25306 MLR #include &lt;sys/mman.h&gt;

25307 int mlock(const void \* *addr*, size\_t *len*);25308 int munlock(const void \* *addr*, size\_t *len*);

25309

25310 **DESCRIPTION**

25311 The *mlock()* function shall cause those whole pages containing any part of the address space of  
 25312 the process starting at address *addr* and continuing for *len* bytes to be memory-resident until  
 25313 unlocked or until the process exits or *execs* another process image. The implementation may  
 25314 require that *addr* be a multiple of {PAGESIZE}.

25315 The *munlock()* function shall unlock those whole pages containing any part of the address space  
 25316 of the process starting at address *addr* and continuing for *len* bytes, regardless of how many  
 25317 times *mlock()* has been called by the process for any of the pages in the specified range. The  
 25318 implementation may require that *addr* be a multiple of the {PAGESIZE}.

25319 If any of the pages in the range specified to a call to *munlock()* are also mapped into the address  
 25320 spaces of other processes, any locks established on those pages by another process are  
 25321 unaffected by the call of this process to *munlock()*. If any of the pages in the range specified by a  
 25322 call to *munlock()* are also mapped into other portions of the address space of the calling process  
 25323 outside the range specified, any locks established on those pages via the other mappings are also  
 25324 unaffected by this call.

25325 Upon successful return from *mlock()*, pages in the specified range shall be locked and memory-  
 25326 resident. Upon successful return from *munlock()*, pages in the specified range shall be unlocked  
 25327 with respect to the address space of the process. Memory residency of unlocked pages is  
 25328 unspecified.

25329 The appropriate privilege is required to lock process memory with *mlock()*.

25330 **RETURN VALUE**

25331 Upon successful completion, the *mlock()* and *munlock()* functions shall return a value of zero.  
 25332 Otherwise, no change is made to any locks in the address space of the process, and the function  
 25333 shall return a value of  $-1$  and set *errno* to indicate the error.

25334 **ERRORS**

25335 The *mlock()* and *munlock()* functions shall fail if:

25336 [ENOMEM] Some or all of the address range specified by the *addr* and *len* arguments does  
 25337 not correspond to valid mapped pages in the address space of the process.

25338 The *mlock()* function shall fail if:

25339 [EAGAIN] Some or all of the memory identified by the operation could not be locked  
 25340 when the call was made.

25341 The *mlock()* and *munlock()* functions may fail if:

25342 [EINVAL] The *addr* argument is not a multiple of {PAGESIZE}.

25343 The *mlock()* function may fail if:

25344 [ENOMEM] Locking the pages mapped by the specified range would exceed an  
 25345 implementation-defined limit on the amount of memory that the process may  
 25346 lock.

25347 [EPERM] The calling process does not have the appropriate privilege to perform the  
25348 requested operation.

25349 **EXAMPLES**

25350 None.

25351 **APPLICATION USAGE**

25352 None.

25353 **RATIONALE**

25354 None.

25355 **FUTURE DIRECTIONS**

25356 None.

25357 **SEE ALSO**

25358 *exec*, *exit()*, *fork()*, *mlockall()*, *munmap()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
25359 <sys/mman.h>

25360 **CHANGE HISTORY**

25361 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25362 **Issue 6**

25363 The *mlock()* and *munlock()* functions are marked as part of the Range Memory Locking option.

25364 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
25365 implementation does not support the Range Memory Locking option.

25366 **NAME**25367 mlockall, munlockall — lock/unlock the address space of a process (**REALTIME**)25368 **SYNOPSIS**25369 ML 

```
#include <sys/mman.h>
```

25370 

```
int mlockall(int flags);
```

25371 

```
int munlockall(void);
```

25372

25373 **DESCRIPTION**

25374 The *mlockall()* function shall cause all of the pages mapped by the address space of a process to  
 25375 be memory-resident until unlocked or until the process exits or *execs* another process image. The  
 25376 *flags* argument determines whether the pages to be locked are those currently mapped by the  
 25377 address space of the process, those that are mapped in the future, or both. The *flags* argument is  
 25378 constructed from the bitwise-inclusive OR of one or more of the following symbolic constants,  
 25379 defined in *<sys/mman.h>*:

25380 MCL\_CURRENT Lock all of the pages currently mapped into the address space of the process.

25381 MCL\_FUTURE Lock all of the pages that become mapped into the address space of the  
25382 process in the future, when those mappings are established.

25383 If MCL\_FUTURE is specified, and the automatic locking of future mappings eventually causes  
 25384 the amount of locked memory to exceed the amount of available physical memory or any other  
 25385 implementation-defined limit, the behavior is implementation-defined. The manner in which the  
 25386 implementation informs the application of these situations is also implementation-defined.

25387 The *munlockall()* function unlocks all currently mapped pages of the address space of the  
 25388 process. Any pages that become mapped into the address space of the process after a call to  
 25389 *munlockall()* shall not be locked, unless there is an intervening call to *mlockall()* specifying  
 25390 MCL\_FUTURE or a subsequent call to *mlockall()* specifying MCL\_CURRENT. If pages mapped  
 25391 into the address space of the process are also mapped into the address spaces of other processes  
 25392 and are locked by those processes, the locks established by the other processes are unaffected by  
 25393 a call by this process to *munlockall()*.

25394 Upon successful return from the *mlockall()* function that specifies MCL\_CURRENT, all currently  
 25395 mapped pages of the process' address space shall be memory-resident and locked. Upon return  
 25396 from the *munlockall()* function, all currently mapped pages of the process' address space shall be  
 25397 unlocked with respect to the process' address space. The memory residency of unlocked pages is  
 25398 unspecified.

25399 The appropriate privilege is required to lock process memory with *mlockall()*.25400 **RETURN VALUE**

25401 Upon successful completion, the *mlockall()* function shall return a value of zero. Otherwise, no  
 25402 additional memory shall be locked, and the function shall return a value of  $-1$  and set *errno* to  
 25403 indicate the error. The effect of failure of *mlockall()* on previously existing locks in the address  
 25404 space is unspecified.

25405 If it is supported by the implementation, the *munlockall()* function shall always return a value of  
 25406 zero. Otherwise, the function shall return a value of  $-1$  and set *errno* to indicate the error.

25407 **ERRORS**25408 The *mlockall()* function shall fail if:

25409 [EAGAIN] Some or all of the memory identified by the operation could not be locked  
 25410 when the call was made.

|       |                          |                                                                                                                                        |  |
|-------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------|--|
| 25411 | [EINVAL]                 | The <i>flags</i> argument is zero, or includes unimplemented flags.                                                                    |  |
| 25412 |                          | The <i>mlockall()</i> function may fail if:                                                                                            |  |
| 25413 | [ENOMEM]                 | Locking all of the pages currently mapped into the address space of the                                                                |  |
| 25414 |                          | process would exceed an implementation-defined limit on the amount of                                                                  |  |
| 25415 |                          | memory that the process may lock.                                                                                                      |  |
| 25416 | [EPERM]                  | The calling process does not have the appropriate privilege to perform the                                                             |  |
| 25417 |                          | requested operation.                                                                                                                   |  |
| 25418 | <b>EXAMPLES</b>          |                                                                                                                                        |  |
| 25419 |                          | None.                                                                                                                                  |  |
| 25420 | <b>APPLICATION USAGE</b> |                                                                                                                                        |  |
| 25421 |                          | None.                                                                                                                                  |  |
| 25422 | <b>RATIONALE</b>         |                                                                                                                                        |  |
| 25423 |                          | None.                                                                                                                                  |  |
| 25424 | <b>FUTURE DIRECTIONS</b> |                                                                                                                                        |  |
| 25425 |                          | None.                                                                                                                                  |  |
| 25426 | <b>SEE ALSO</b>          |                                                                                                                                        |  |
| 25427 |                          | <i>exec</i> , <i>exit()</i> , <i>fork()</i> , <i>mlock()</i> , <i>munmap()</i> , the Base Definitions volume of IEEE Std. 1003.1-200x, |  |
| 25428 |                          | < <i>sys/mman.h</i> >                                                                                                                  |  |
| 25429 | <b>CHANGE HISTORY</b>    |                                                                                                                                        |  |
| 25430 |                          | First released in Issue 5. Included for alignment with the POSIX Realtime Extension.                                                   |  |
| 25431 | <b>Issue 6</b>           |                                                                                                                                        |  |
| 25432 |                          | The <i>mlockall()</i> and <i>munlockall()</i> functions are marked as part of the Process Memory Locking                               |  |
| 25433 |                          | option.                                                                                                                                |  |
| 25434 |                          | The [ENOSYS] error condition has been removed as stubs need not be provided if an                                                      |  |
| 25435 |                          | implementation does not support the Process Memory Locking option.                                                                     |  |

## 25436 NAME

25437 mmap — map pages of memory

## 25438 SYNOPSIS

25439 MF|SHM #include &lt;sys/mman.h&gt;

```
25440 void *mmap(void *addr, size_t len, int prot, int flags,
25441 int fildes, off_t off);
25442
```

## 25443 DESCRIPTION

25444 The *mmap()* function shall establish a mapping between a process' address space and a file, shared memory object, or typed memory object. The format of the call is as follows:

```
25446 pa=mmap(addr, len, prot, flags, fildes, off);
```

25447 The *mmap()* function establishes a mapping between the address space of the process at an address *pa* for *len* bytes to the memory object represented by the file descriptor *fildes* at offset *off* for *len* bytes. The value of *pa* is an implementation-defined function of the parameter *addr* and the values of *flags*, further described below. A successful *mmap()* call shall return *pa* as its result. The address range starting at *pa* and continuing for *len* bytes shall be legitimate for the possible (not necessarily current) address space of the process. The range of bytes starting at *off* and continuing for *len* bytes shall be legitimate for the possible (not necessarily current) offsets in the file, shared memory object, or typed memory object represented by *fildes*.

25455 TYM If *fildes* represents a typed memory object opened with either the POSIX\_TYPED\_MEM\_ALLOCATE flag or the POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG flag, the memory object to be mapped shall be that portion of the typed memory object allocated by the implementation as specified below. In this case, if *off* is non-zero, the behavior of *mmap()* is undefined. If *fildes* refers to a valid typed memory object that is not accessible from the calling process, *mmap()* shall fail.

25461 The mapping established by *mmap()* replaces any previous mappings for those whole pages containing any part of the address space of the process starting at *pa* and continuing for *len* bytes.

25464 If the size of the mapped file changes after the call to *mmap()* as a result of some other operation on the mapped file, the effect of references to portions of the mapped region that correspond to added or removed portions of the file is unspecified.

25467 TYM The *mmap()* function is supported for regular files, shared memory objects, and typed memory objects. Support for any other type of file is unspecified.

25469 The parameter *prot* determines whether read, write, execute, or some combination of accesses are permitted to the data being mapped. The *prot* should be either PROT\_NONE or the bitwise-inclusive OR of one or more of the other flags in the following table, defined in the header <sys/mman.h>.

25473

25474

25475

25476

25477

25478

| Symbolic Constant | Description              |
|-------------------|--------------------------|
| PROT_READ         | Data can be read.        |
| PROT_WRITE        | Data can be written.     |
| PROT_EXEC         | Data can be executed.    |
| PROT_NONE         | Data cannot be accessed. |

25479 If an implementation cannot support the combination of access types specified by *prot*, the call to *mmap()* fails. An implementation may permit accesses other than those specified by *prot*; however, if the Memory Protection option is supported, the implementation shall not permit a

25480 MPR

25481

25482 write to succeed where PROT\_WRITE has not been set or shall not permit any access where  
 25483 PROT\_NONE alone has been set. The implementation shall support at least the following values  
 25484 of *prot*: PROT\_NONE, PROT\_READ, PROT\_WRITE, and the bitwise-inclusive OR of  
 25485 PROT\_READ and PROT\_WRITE. If the Memory Protection option is not supported, the result of  
 25486 any access that conflicts with the specified protection is undefined. The file descriptor *fdes* shall  
 25487 have been opened with read permission, regardless of the protection options specified. If  
 25488 PROT\_WRITE is specified, the application shall ensure that it has opened the file descriptor  
 25489 *fdes* with write permission unless MAP\_PRIVATE is specified in the *flags* parameter as  
 25490 described below.

25491 The parameter *flags* provides other information about the handling of the mapped data. The  
 25492 value of *flags* is the bitwise-inclusive OR of these options, defined in <sys/mman.h>:

25493

25494

25495

25496

25497

| Symbolic Constant | Description                    |
|-------------------|--------------------------------|
| MAP_SHARED        | Changes are shared.            |
| MAP_PRIVATE       | Changes are private.           |
| MAP_FIXED         | Interpret <i>addr</i> exactly. |

25498 Implementations that do not support the Memory Mapped Files option are not required to  
 25499 XSI support MAP\_PRIVATE. It is implementation-defined whether MAP\_FIXED shall be supported.

25500 MAP\_FIXED shall be supported on XSI-conformant systems.

25501 MAP\_SHARED and MAP\_PRIVATE describe the disposition of write references to the memory  
 25502 object. If MAP\_SHARED is specified, write references change the underlying object. If  
 25503 MAP\_PRIVATE is specified, modifications to the mapped data by the calling process shall be  
 25504 visible only to the calling process and shall not change the underlying object. It is unspecified  
 25505 whether modifications to the underlying object done after the MAP\_PRIVATE mapping is  
 25506 established are visible through the MAP\_PRIVATE mapping. Either MAP\_SHARED or  
 25507 MAP\_PRIVATE can be specified, but not both. The mapping type is retained across *fork*().

25508 TYM When *fdes* represents a typed memory object opened with either the  
 25509 POSIX\_TYPED\_MEM\_ALLOCATE flag or the POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG  
 25510 flag, *mmap*( ) shall, if there are enough resources available, map *len* bytes allocated from the  
 25511 corresponding typed memory object which were not previously allocated to any process in any  
 25512 processor that may access that typed memory object. If there are not enough resources available,  
 25513 the function shall fail. If *fdes* represents a typed memory object opened with the  
 25514 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG flag, these allocated bytes shall be contiguous  
 25515 within the typed memory object. If *fdes* represents a typed memory object opened with the  
 25516 POSIX\_TYPED\_MEM\_ALLOCATE flag, these allocated bytes may be composed of non-  
 25517 contiguous fragments within the typed memory object. If *fdes* represents a typed memory  
 25518 object opened with neither the POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG flag nor the  
 25519 POSIX\_TYPED\_MEM\_ALLOCATE flag, *len* bytes starting at offset *off* within the typed memory  
 25520 object are mapped, exactly as when mapping a file or shared memory object. In this case, if two  
 25521 processes map an area of typed memory using the same *off* and *len* values and using file  
 25522 descriptors that refer to the same memory pool (either from the same port or from a different  
 25523 port), both processes shall map the same region of storage.

25524 When MAP\_FIXED is set in the *flags* argument, the implementation is informed that the value of  
 25525 *pa* shall be *addr*, exactly. If MAP\_FIXED is set, *mmap*( ) may return MAP\_FAILED and set *errno* to  
 25526 [EINVAL]. If a MAP\_FIXED request is successful, the mapping established by *mmap*( ) replaces  
 25527 any previous mappings for the process' pages in the range [*pa*,*pa+len*).

25528 When MAP\_FIXED is not set, the implementation uses *addr* in an implementation-defined  
 25529 manner to arrive at *pa*. The *pa* so chosen shall be an area of the address space that the

- 25530 implementation deems suitable for a mapping of *len* bytes to the file. All implementations  
 25531 interpret an *addr* value of 0 as granting the implementation complete freedom in selecting *pa*,  
 25532 subject to constraints described below. A non-zero value of *addr* is taken to be a suggestion of a  
 25533 process address near which the mapping should be placed. When the implementation selects a  
 25534 value for *pa*, it never places a mapping at address 0, nor does it replace any extant mapping.
- 25535 The *off* argument is constrained to be aligned and sized according to the value returned by  
 25536 *sysconf()* when passed *\_SC\_PAGESIZE* or *\_SC\_PAGE\_SIZE*. When *MAP\_FIXED* is specified, the  
 25537 application shall ensure that the argument *addr* also meets these constraints. The  
 25538 implementation performs mapping operations over whole pages. Thus, while the argument *len*  
 25539 need not meet a size or alignment constraint, the implementation shall include, in any mapping  
 25540 operation, any partial page specified by the range [*pa*,*pa+len*).
- 25541 The system shall always zero-fill any partial page at the end of an object. Further, the system  
 25542 shall never write out any modified portions of the last page of an object which are beyond its  
 25543 MPR end. References within the address range starting at *pa* and continuing for *len* bytes to whole  
 25544 pages following the end of an object shall result in delivery of a SIGBUS signal.
- 25545 An implementation may generate SIGBUS signals when a reference would cause an error in the  
 25546 mapped object, such as out-of-space condition.
- 25547 The *mmap()* function adds an extra reference to the file associated with the file descriptor *fildes*  
 25548 which is not removed by a subsequent *close()* on that file descriptor. This reference is removed  
 25549 when there are no more mappings to the file.
- 25550 The *st\_atime* field of the mapped file may be marked for update at any time between the *mmap()*  
 25551 call and the corresponding *munmap()* call. The initial read or write reference to a mapped region  
 25552 shall cause the file's *st\_atime* field to be marked for update if it has not already been marked for  
 25553 update.
- 25554 The *st\_ctime* and *st\_mtime* fields of a file that is mapped with *MAP\_SHARED* and *PROT\_WRITE*  
 25555 shall be marked for update at some point in the interval between a write reference to the  
 25556 mapped region and the next call to *msync()* with *MS\_ASYNC* or *MS\_SYNC* for that portion of  
 25557 the file by any process. If there is no such call and if the underlying file is modified as a result of  
 25558 a write reference, then these fields shall be marked for update at some time after the write  
 25559 reference.
- 25560 There may be implementation-defined limits on the number of memory regions that can be  
 25561 mapped (per process or per system).
- 25562 XSI If such a limit is imposed, whether the number of memory regions that can be mapped by a  
 25563 process is decreased by the use of *shmat()* is implementation-defined.
- 25564 If *mmap()* fails for reasons other than [EBADF], [EINVAL], or [ENOTSUP], some of the  
 25565 mappings in the address range starting at *addr* and continuing for *len* bytes may have been  
 25566 unmapped.
- 25567 **RETURN VALUE**
- 25568 Upon successful completion, the *mmap()* function shall return the address at which the mapping  
 25569 was placed (*pa*); otherwise, it shall return a value of *MAP\_FAILED* and set *errno* to indicate the  
 25570 error. The symbol *MAP\_FAILED* is defined in the header <sys/mman.h>. No successful return  
 25571 from *mmap()* shall return the value *MAP\_FAILED*.
- 25572 **ERRORS**
- 25573 The *mmap()* function shall fail if:
- 25574 [EACCES] The *fildes* argument is not open for read, regardless of the protection specified,  
 25575 or *fildes* is not open for write and *PROT\_WRITE* was specified for a

|                         |             |                                                                                                                                                                                                                                        |
|-------------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 25576                   |             | MAP_SHARED type mapping.                                                                                                                                                                                                               |
| 25577 ML<br>25578       | [EAGAIN]    | The mapping could not be locked in memory, if required by <i>mlockall()</i> , due to a lack of resources.                                                                                                                              |
| 25579                   | [EBADF]     | The <i>fildev</i> argument is not a valid open file descriptor.                                                                                                                                                                        |
| 25580<br>25581<br>25582 | [EINVAL]    | The <i>addr</i> argument (if MAP_FIXED was specified) or <i>off</i> is not a multiple of the page size as returned by <i>sysconf()</i> , or are considered invalid by the implementation.                                              |
| 25583<br>25584          | [EINVAL]    | The value of <i>flags</i> is invalid (neither MAP_PRIVATE nor MAP_SHARED is set).                                                                                                                                                      |
| 25585<br>25586          | [EMFILE]    | The number of mapped regions would exceed an implementation-defined limit (per process or per system).                                                                                                                                 |
| 25587                   | [ENODEV]    | The <i>fildev</i> argument refers to a file whose type is not supported by <i>mmap()</i> .                                                                                                                                             |
| 25588<br>25589<br>25590 | [ENOMEM]    | MAP_FIXED was specified, and the range [ <i>addr,addr+len</i> ) exceeds that allowed for the address space of a process; or, if MAP_FIXED was not specified and there is insufficient room in the address space to effect the mapping. |
| 25591 ML<br>25592       | [ENOMEM]    | The mapping could not be locked in memory, if required by <i>mlockall()</i> , because it would require more space than the system is able to supply.                                                                                   |
| 25593<br>25594          |             | MAP_FIXED or MAP_PRIVATE was specified in the <i>flags</i> argument and the implementation does not support this functionality.                                                                                                        |
| 25595 TYM<br>25596      | [ENOMEM]    | Not enough unallocated memory resources remain in the typed memory object designated by <i>fildev</i> to allocate <i>len</i> bytes.                                                                                                    |
| 25597<br>25598          | [ENOTSUP]   | The implementation does not support the combination of accesses requested in the <i>prot</i> argument.                                                                                                                                 |
| 25599                   | [ENXIO]     | Addresses in the range [ <i>off,off+len</i> ) are invalid for the object specified by <i>fildev</i> .                                                                                                                                  |
| 25600<br>25601          | [ENXIO]     | MAP_FIXED was specified in <i>flags</i> and the combination of <i>addr</i> , <i>len</i> , and <i>off</i> is invalid for the object specified by <i>fildev</i> .                                                                        |
| 25602 TYM<br>25603      | [ENXIO]     | The <i>fildev</i> argument refers to a typed memory object that is not accessible from the calling process.                                                                                                                            |
| 25604<br>25605          | [EOVERFLOW] | The file is a regular file and the value of <i>off</i> plus <i>len</i> exceeds the offset maximum established in the open file description associated with <i>fildev</i> .                                                             |

#### 25606 EXAMPLES

25607 None.

#### 25608 APPLICATION USAGE

25609 Use of *mmap()* may reduce the amount of memory available to other memory allocation  
25610 functions.

25611 Use of MAP\_FIXED may result in unspecified behavior in further use of *malloc()* and *shmat()*.  
25612 The use of MAP\_FIXED is discouraged, as it may prevent an implementation from making the  
25613 most effective use of resources.

25614 The application must ensure correct synchronization when using *mmap()* in conjunction with  
25615 any other file access method, such as *read()* and *write()*, standard input/output, and *shmat()*.

25616 The *mmap()* function allows access to resources via address space manipulations, instead of  
25617 *read()/write()*. Once a file is mapped, all a process has to do to access it is use the data at the

25618 address to which the file was mapped. So, using pseudo-code to illustrate the way in which an  
 25619 existing program might be changed to use *mmap()*, the following:

```
25620 fildes = open(...)
25621 lseek(fildes, some_offset)
25622 read(fildes, buf, len)
25623 /* Use data in buf. */
```

25624 becomes:

```
25625 fildes = open(...)
25626 address = mmap(0, len, PROT_READ, MAP_PRIVATE, fildes, some_offset)
25627 /* Use data at address. */
```

25628 The [EINVAL] error above is marked EX because it is defined as an optional error in the POSIX  
 25629 Realtime Extension.

### 25630 RATIONALE

25631 After considering several other alternatives, it was decided to adopt the *mmap()* definition found  
 25632 in SVR4 for mapping memory objects into process address spaces. The SVR4 definition is  
 25633 minimal, in that it describes only what has been built, and what appears to be necessary for a  
 25634 general and portable mapping facility.

25635 Note that while *mmap()* was first designed for mapping files, it is actually a general-purpose  
 25636 mapping facility. It can be used to map any appropriate object, such as memory, files, devices,  
 25637 and so on, into the address space of a process.

25638 When a mapping is established, it is possible that the implementation may need to map more  
 25639 than is requested into the address space of the process because of hardware requirements. An  
 25640 application, however, cannot count on this behavior. Implementations that do not use a paged  
 25641 architecture may simply allocate a common memory region and return the address of it; such  
 25642 implementations probably do not allocate any more than is necessary. References past the end of  
 25643 the requested area are unspecified.

25644 If an application requests a mapping that would overlay existing mappings in the process, it  
 25645 might be desirable that an implementation detect this and inform the application. However, the  
 25646 default, portable (not MAP\_FIXED) operation does not overlay existing mappings. On the other  
 25647 hand, if the program specifies a fixed address mapping (which requires some implementation  
 25648 knowledge to determine a suitable address, if the function is supported at all), then the program  
 25649 is presumed to be successfully managing its own address space and should be trusted when it  
 25650 asks to map over existing data structures. Furthermore, it is also desirable to make as few system  
 25651 calls as possible, and it might be considered onerous to require an *munmap()* before an *mmap()*  
 25652 to the same address range. This volume of IEEE Std. 1003.1-200x specifies that the new  
 25653 mappings replace any existing mappings, following existing practice in this regard.

25654 It is not expected, when the Memory Protection option is supported, that all hardware  
 25655 implementations are able to support all combinations of permissions at all addresses. When this  
 25656 option is supported, implementations are required to disallow write access to mappings without  
 25657 write permission and to disallow access to mappings without any access permission. Other than  
 25658 these restrictions, implementations may allow access types other than those requested by the  
 25659 application. For example, if the application requests only PROT\_WRITE, the implementation  
 25660 may also allow read access. A call to *mmap()* fails if the implementation cannot support allowing  
 25661 all the access requested by the application. For example, some implementations cannot support  
 25662 a request for both write access and execute access simultaneously. All implementations  
 25663 supporting the Memory Protection option must support requests for no access, read access,  
 25664 write access, and both read and write access. Strictly conforming code must only rely on the  
 25665 required checks. These restrictions allow for portability across a wide range of hardware.

25666 The MAP\_FIXED address treatment is likely to fail for non-page-aligned values and for certain  
25667 architecture-dependent address ranges. Conforming implementations cannot count on being  
25668 able to choose address values for MAP\_FIXED without utilizing non-portable, implementation-  
25669 defined knowledge. Nonetheless, MAP\_FIXED is provided as a standard interface conforming to  
25670 existing practice for utilizing such knowledge when it is available.

25671 Similarly, in order to allow implementations that do not support virtual addresses, support for  
25672 directly specifying any mapping addresses via MAP\_FIXED is not required and thus a  
25673 conforming application may not count on it.

25674 The MAP\_PRIVATE function can be implemented efficiently when memory protection hardware  
25675 is available. When such hardware is not available, implementations can implement such  
25676 “mappings” by simply making a real copy of the relevant data into process private memory,  
25677 though this tends to behave similarly to *read()*.

25678 The function has been defined to allow for many different models of using shared memory.  
25679 However, all uses are not equally portable across all machine architectures. In particular, the  
25680 *mmap()* function allows the system as well as the application to specify the address at which to  
25681 map a specific region of a memory object. The most portable way to use the function is always to  
25682 let the system choose the address, specifying NULL as the value for the argument *addr* and not  
25683 to specify MAP\_FIXED.

25684 If it is intended that a particular region of a memory object be mapped at the same address in a  
25685 group of processes (on machines where this is even possible), then MAP\_FIXED can be used to  
25686 pass in the desired mapping address. The system can still be used to choose the desired address  
25687 if the first such mapping is made without specifying MAP\_FIXED, and then the resulting  
25688 mapping address can be passed to subsequent processes for them to pass in via MAP\_FIXED.  
25689 The availability of a specific address range cannot be guaranteed, in general.

25690 The *mmap()* function can be used to map a region of memory that is larger than the current size  
25691 of the object. Memory access within the mapping but beyond the current end of the underlying  
25692 objects may result in SIGBUS signals being sent to the process. The reason for this is that the size  
25693 of the object can be manipulated by other processes and can change at any moment. The  
25694 implementation should tell the application that a memory reference is outside the object where  
25695 this can be detected; otherwise, written data may be lost and read data may not reflect actual  
25696 data in the object.

25697 Note that references beyond the end of the object do not extend the object as the new end cannot  
25698 be determined precisely by most virtual memory hardware. Instead, the size can be directly  
25699 manipulated by *ftruncate()*.

25700 Process memory locking does apply to shared memory regions, and the MEMLOCK\_FUTURE  
25701 argument to *memlockall()* can be relied upon to cause new shared memory regions to be  
25702 automatically locked.

25703 Existing implementations of *mmap()* return the value `-1` when unsuccessful. Since the casting of  
25704 this value to type `void*` cannot be guaranteed by the ISO C standard to be distinct from a  
25705 successful value, this volume of IEEE Std. 1003.1-200x defines the symbol MAP\_FAILED, which  
25706 a conforming implementation does not return as the result of a successful call.

#### 25707 FUTURE DIRECTIONS

25708 None.

#### 25709 SEE ALSO

25710 *exec*, *fcntl()*, *fork()*, *lockf()*, *msync()*, *munmap()*, *mprotect()*, *posix\_typed\_mem\_open()*, *shmat()*,  
25711 *sysconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/mman.h>

25712 **CHANGE HISTORY**

25713 First released in Issue 4, Version 2.

25714 **Issue 5**

25715 Moved from X/OPEN UNIX extension to BASE.

25716 Aligned with *mmap()* in the POSIX Realtime Extension as follows:

- 25717 • The DESCRIPTION is extensively reworded.
- 25718 • The [EAGAIN] and [ENOTSUP] mandatory error conditions are added.
- 25719 • New cases of [ENOMEM] and [ENXIO] are added as mandatory error conditions.
- 25720 • The value returned on failure is the value of the constant MAP\_FAILED; this was previously
- 25721 defined as -1.

25722 Large File Summit extensions are added.

25723 **Issue 6**25724 The *mmap()* function is marked as part of the Memory Mapped Files option.25725 The Open Group corrigenda item U028/6 has been applied, changing (void \*)-1 to  
25726 MAP\_FAILED.25727 The following new requirements on POSIX implementations derive from alignment with the  
25728 Single UNIX Specification:

- 25729 • The DESCRIPTION is updated to described the use of MAP\_FIXED.
- 25730 • The DESCRIPTION is updated to describe the addition of an extra reference to the file
- 25731 associated with the file descriptor passed to *mmap()*.
- 25732 • The DESCRIPTION is updated to state that there may be implementation-defined limits on
- 25733 the number of memory regions that can be mapped.
- 25734 • The DESCRIPTION is updated to describe constraints on the alignment and size of the *off*
- 25735 argument.
- 25736 • The [EINVAL] and [EMFILE] error conditions are added.
- 25737 • The [EOVERFLOW] error condition is added. This change is to support large files.

25738 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 25739 • The DESCRIPTION is updated to describe the cases when MAP\_PRIVATE and MAP\_FIXED
- 25740 need not be supported.

25741 The following changes are made for alignment with IEEE Std. 1003.1j-2000:

- 25742 • Semantics for typed memory objects are added to the DESCRIPTION.
- 25743 • New [ENOMEM] and [ENXIO] errors are added to the ERRORS section.
- 25744 • The *posix\_typed\_mem\_open()* function is added to the SEE ALSO section.

25745 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

25746 **NAME**

25747 modf, modff, modfl — decompose a floating-point number

25748 **SYNOPSIS**

25749 #include <math.h>

25750 double modf(double *x*, double \**iptr*);

25751 float modff(float *value*, float \**iptr*);

25752 long double modfl(long double *value*, long double \**iptr*);

25753 **DESCRIPTION**

25754 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
25755 conflict between the requirements described here and the ISO C standard is unintentional. This  
25756 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

25757 These functions shall break the argument *x* into integral and fractional parts, each of which has  
25758 the same sign as the argument. It stores the integral part as a double in the object pointed to by  
25759 *iptr*.

25760 An application wishing to check for error situations should set *errno* to 0 before calling *modf()*. If  
25761 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

25762 **RETURN VALUE**

25763 Upon successful completion, these functions shall return the signed fractional part of *x*.

25764 XSI If *x* is NaN, NaN shall be returned, *errno* may be set to [EDOM], and \**iptr* shall be set to NaN.

25765 If the correct value would cause underflow, 0 shall be returned and *errno* may be set to  
25766 [ERANGE].

25767 **ERRORS**

25768 These functions may fail if:

25769 XSI [EDOM] The value of *x* is NaN.

25770 [ERANGE] The result underflows.

25771 XSI No other errors shall occur.

25772 **EXAMPLES**

25773 None.

25774 **APPLICATION USAGE**

25775 The *modf()* function computes the function result and \**iptr* such that:

25776 a = modf(x, &iptr)

25777 x == a+iptr

25778 allowing for the usual floating point inaccuracies.

25779 **RATIONALE**

25780 None.

25781 **FUTURE DIRECTIONS**

25782 None.

25783 **SEE ALSO**

25784 *frexp()*, *isnan()*, *ldexp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

25785 **CHANGE HISTORY**

25786 First released in Issue 1. Derived from Issue 1 of the SVID.

25787 **Issue 4**

25788 References to *matherr()* are removed.

25789 The name of the first argument is changed from *value* to *x*.

25790 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the ISO C standard and to rationalize error handling in the mathematics functions.

25792 The return value specified for [EDOM] is marked as an extension.

25793 **Issue 5**

25794 The DESCRIPTION is updated to indicate how an application should check for an error. This text was previously published in the APPLICATION USAGE section.

25796 **Issue 6**

25797 The *modff()* and *modfl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

25799 **NAME**

25800 mprotect — set protection of memory mapping

25801 **SYNOPSIS**

25802 MPR #include &lt;sys/mman.h&gt;

25803 int mprotect(void \*addr, size\_t len, int prot);

25804

25805 **DESCRIPTION**

25806 The *mprotect()* function shall change the access protections to be that specified by *prot* for those  
 25807 whole pages containing any part of the address space of the process starting at address *addr* and  
 25808 continuing for *len* bytes. The parameter *prot* determines whether read, write, execute, or some  
 25809 combination of accesses are permitted to the data being mapped. The *prot* argument should be  
 25810 either PROT\_NONE or the bitwise-inclusive OR of one or more of PROT\_READ, PROT\_WRITE,  
 25811 and PROT\_EXEC.

25812 If an implementation cannot support the combination of access types specified by *prot*, the call  
 25813 to *mprotect()* shall fail.

25814 An implementation may permit accesses other than those specified by *prot*; however, no  
 25815 implementation shall permit a write to succeed where PROT\_WRITE has not been set or shall  
 25816 permit any access where PROT\_NONE alone has been set. Implementations shall support at  
 25817 least the following values of *prot*: PROT\_NONE, PROT\_READ, PROT\_WRITE, and the bitwise-  
 25818 inclusive OR of PROT\_READ and PROT\_WRITE. If PROT\_WRITE is specified, the application  
 25819 shall ensure that it has opened the mapped objects in the specified address range with write  
 25820 permission, unless MAP\_PRIVATE was specified in the original mapping, regardless of whether  
 25821 the file descriptors used to map the objects have since been closed.

25822 The implementation shall require that *addr* be a multiple of the page size as returned by  
 25823 *sysconf()*.

25824 The behavior of this function is unspecified if the mapping was not established by a call to  
 25825 *mmap()*.

25826 When *mprotect()* fails for reasons other than [EINVAL], the protections on some of the pages in  
 25827 the range [*addr*,*addr*+*len*) may have been changed.

25828 **RETURN VALUE**

25829 Upon successful completion, *mprotect()* shall return 0; otherwise, it shall return -1 and set *errno*  
 25830 to indicate the error.

25831 **ERRORS**25832 The *mprotect()* function shall fail if:

25833 [EACCES] The *prot* argument specifies a protection that violates the access permission  
 25834 the process has to the underlying memory object.

25835 [EAGAIN] The *prot* argument specifies PROT\_WRITE over a MAP\_PRIVATE mapping  
 25836 and there are insufficient memory resources to reserve for locking the private  
 25837 page.

25838 [EINVAL] The *addr* argument is not a multiple of the page size as returned by *sysconf()*.

25839 [ENOMEM] Addresses in the range [*addr*,*addr*+*len*) are invalid for the address space of a  
 25840 process, or specify one or more pages which are not mapped.

25841 [ENOMEM] The *prot* argument specifies PROT\_WRITE on a MAP\_PRIVATE mapping, and  
 25842 it would require more space than the system is able to supply for locking the  
 25843 private pages, if required.

25844 [ENOTSUP] The implementation does not support the combination of accesses requested  
25845 in the *prot* argument.

25846 **EXAMPLES**

25847 None.

25848 **APPLICATION USAGE**

25849 The [EINVAL] error above is marked EX because it is defined as an optional error in the POSIX  
25850 Realtime Extension.

25851 **RATIONALE**

25852 None.

25853 **FUTURE DIRECTIONS**

25854 None.

25855 **SEE ALSO**

25856 *mmap()*, *sysconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/mman.h>

25857 **CHANGE HISTORY**

25858 First released in Issue 4, Version 2.

25859 **Issue 5**

25860 Moved from X/OPEN UNIX extension to BASE.

25861 Aligned with *mprotect()* in the POSIX Realtime Extension as follows:

- 25862 • The DESCRIPTION is largely reworded.
- 25863 • [ENOTSUP] and a second form of [ENOMEM] are added as mandatory error conditions.
- 25864 • [EAGAIN] is moved from the optional to the mandatory error conditions.

25865 **Issue 6**

25866 The *mprotect()* function is marked as part of the Memory Protection option.

25867 The following new requirements on POSIX implementations derive from alignment with the  
25868 Single UNIX Specification:

- 25869 • The DESCRIPTION is updated to state that implementations require *addr* to be a multiple of  
25870 the page size as returned by *sysconf()*.
- 25871 • The [EINVAL] error condition is added.

25872 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

25873 **NAME**25874 mq\_close — close a message queue (**REALTIME**)25875 **SYNOPSIS**

25876 MSG #include &lt;mqueue.h&gt;

25877 int mq\_close(mqd\_t mqdes);

25878

25879 **DESCRIPTION**

25880 The *mq\_close()* function shall remove the association between the message queue descriptor,  
25881 *mqdes*, and its message queue. The results of using this message queue descriptor after  
25882 successful return from this *mq\_close()*, and until the return of this message queue descriptor  
25883 from a subsequent *mq\_open()*, are undefined.

25884 If the process has successfully attached a notification request to the message queue via this  
25885 *mqdes*, this attachment shall be removed, and the message queue is available for another process  
25886 to attach for notification.

25887 **RETURN VALUE**

25888 Upon successful completion, the *mq\_close()* function shall return a value of zero; otherwise, the  
25889 function shall return a value of -1 and set *errno* to indicate the error.

25890 **ERRORS**25891 The *mq\_close()* function shall fail if:

25892 [EBADF] The *mqdes* argument is not a valid message queue descriptor.

25893 **EXAMPLES**

25894 None.

25895 **APPLICATION USAGE**

25896 None.

25897 **RATIONALE**

25898 None.

25899 **FUTURE DIRECTIONS**

25900 None.

25901 **SEE ALSO**

25902 *mq\_open()*, *mq\_unlink()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of  
25903 IEEE Std. 1003.1-200x, <mqueue.h>

25904 **CHANGE HISTORY**

25905 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25906 **Issue 6**25907 The *mq\_close()* function is marked as part of the Message Passing option.

25908 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
25909 implementation does not support the Message Passing option.

25910 **NAME**25911 mq\_getattr — get message queue attributes (**REALTIME**)25912 **SYNOPSIS**

25913 MSG #include &lt;mqqueue.h&gt;

25914 int mq\_getattr(mqd\_t mqdes, struct mq\_attr \*mqstat);

25915

25916 **DESCRIPTION**25917 The *mqdes* argument specifies a message queue descriptor.25918 The *mq\_getattr()* function is used to get status information and attributes of the message queue and the open message queue description associated with the message queue descriptor.25920 The results are returned in the **mq\_attr** structure referenced by the *mqstat* argument.25921 Upon return, the following members have the values associated with the open message queue description as set when the message queue was opened and as modified by subsequent *mq\_setattr()* calls: *mq\_flags*.25924 The following attributes of the message queue shall be returned as set at message queue creation: *mq\_maxmsg*, *mq\_msgsize*.25926 Upon return, the following members within the **mq\_attr** structure referenced by the *mqstat* argument are set to the current state of the message queue:25928 *mq\_curmsgs* The number of messages currently on the queue.25929 **RETURN VALUE**25930 Upon successful completion, the *mq\_getattr()* function shall return zero. Otherwise, the function shall return  $-1$  and set *errno* to indicate the error.25932 **ERRORS**25933 The *mq\_getattr()* function shall fail if:25934 [EBADF] The *mqdes* argument is not a valid message queue descriptor.25935 **EXAMPLES**

25936 None.

25937 **APPLICATION USAGE**

25938 None.

25939 **RATIONALE**

25940 None.

25941 **FUTURE DIRECTIONS**

25942 None.

25943 **SEE ALSO**25944 *mq\_open()*, *mq\_send()*, *mq\_setattr()*, *mq\_timedsend()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**mqqueue.h**>25946 **CHANGE HISTORY**

25947 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25948 **Issue 6**25949 The *mq\_getattr()* function is marked as part of the Message Passing option.

25950 The [ENOSYS] error condition has been removed as stubs need not be provided if an implementation does not support the Message Passing option.

25951

25952           The *mq\_timedsend()* function is added to the SEE ALSO section for alignment with  
25953           IEEE Std. 1003.1d-1999.

25954 **NAME**25955 mq\_notify — notify process that a message is available (**REALTIME**)25956 **SYNOPSIS**

25957 MSG #include &lt;mqueue.h&gt;

25958 int mq\_notify(mqd\_t mqdes, const struct sigevent \*notification);

25959

25960 **DESCRIPTION**

25961 If the argument *notification* is not NULL, this function registers the calling process to be notified  
 25962 of message arrival at an empty message queue associated with the specified message queue  
 25963 descriptor, *mqdes*. The notification specified by the *notification* argument shall be sent to the  
 25964 process when the message queue transitions from empty to non-empty. At any time, only one  
 25965 process may be registered for notification by a message queue. If the calling process or any other  
 25966 process has already registered for notification of message arrival at the specified message queue,  
 25967 subsequent attempts to register for that message queue fail.

25968 If *notification* is NULL and the process is currently registered for notification by the specified  
 25969 message queue, the existing registration is removed.

25970 When the notification is sent to the registered process, its registration shall be removed. The  
 25971 message queue shall then be available for registration.

25972 If a process has registered for notification of message arrival at a message queue and some  
 25973 thread is blocked in *mq\_receive()* waiting to receive a message when a message arrives at the  
 25974 queue, the arriving message satisfies the appropriate *mq\_receive()*. The resulting behavior is as if  
 25975 the message queue remains empty, and no notification is sent.

25976 **RETURN VALUE**

25977 Upon successful completion, the *mq\_notify()* function shall return a value of zero; otherwise, the  
 25978 function shall return a value of -1 and set *errno* to indicate the error.

25979 **ERRORS**25980 The *mq\_notify()* function shall fail if:25981 [EBADF] The *mqdes* argument is not a valid message queue descriptor. |

25982 [EBUSY] A process is already registered for notification by the message queue. |

25983 **EXAMPLES**

25984 None.

25985 **APPLICATION USAGE**

25986 None.

25987 **RATIONALE**

25988 None.

25989 **FUTURE DIRECTIONS**

25990 None.

25991 **SEE ALSO**

25992 *mq\_open()*, *mq\_send()*, *mq\_timedsend()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base  
 25993 Definitions volume of IEEE Std. 1003.1-200x, <mqueue.h>

25994 **CHANGE HISTORY**

25995 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

25996 **Issue 6**

25997 The *mq\_notify()* function is marked as part of the Message Passing option.

25998 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
25999 implementation does not support the Message Passing option.

26000 The *mq\_timedsend()* function is added to the SEE ALSO section for alignment with  
26001 IEEE Std. 1003.1d-1999.

## 26002 NAME

26003 mq\_open — open a message queue (**REALTIME**)

## 26004 SYNOPSIS

26005 MSG #include &lt;mqueue.h&gt;

26006 mqd\_t mq\_open(const char \*name, int oflag, ...);

26007

## 26008 DESCRIPTION

26009 The *mq\_open()* function shall establish the connection between a process and a message queue  
 26010 with a message queue descriptor. It creates an open message queue description that refers to the  
 26011 message queue, and a message queue descriptor that refers to that open message queue  
 26012 description. The message queue descriptor is used by other functions to refer to that message  
 26013 queue. The *name* argument points to a string naming a message queue. It is unspecified whether  
 26014 the name appears in the file system and is visible to other functions that take path names as  
 26015 arguments. The *name* argument conforms to the construction rules for a path name. If *name*  
 26016 begins with the slash character, then processes calling *mq\_open()* with the same value of *name*  
 26017 refer to the same message queue object, as long as that name has not been removed. If *name* does  
 26018 not begin with the slash character, the effect is implementation-defined. The interpretation of  
 26019 slash characters other than the leading slash character in *name* is implementation-defined. If the  
 26020 *name* argument is not the name of an existing message queue and creation is not requested,  
 26021 *mq\_open()* shall fail and return an error.

26022 A message queue descriptor may be implemented using a file descriptor, in which case  
 26023 applications can open up to at least {OPEN\_MAX} file and message queues.

26024 The *oflag* argument requests the desired receive and/or send access to the message queue. The  
 26025 requested access permission to receive messages or send messages is granted if the calling  
 26026 process would be granted read or write access, respectively, to an equivalently protected file.

26027 The value of *oflag* is the bitwise-inclusive OR of values from the following list. Applications  
 26028 specify exactly one of the first three values (access modes) below in the value of *oflag*:

26029 **O\_RDONLY** Open the message queue for receiving messages. The process can use the  
 26030 returned message queue descriptor with *mq\_receive()*, but not *mq\_send()*. A  
 26031 message queue may be open multiple times in the same or different processes  
 26032 for receiving messages.

26033 **O\_WRONLY** Open the queue for sending messages. The process can use the returned  
 26034 message queue descriptor with *mq\_send()* but not *mq\_receive()*. A message  
 26035 queue may be open multiple times in the same or different processes for  
 26036 sending messages.

26037 **O\_RDWR** Open the queue for both receiving and sending messages. The process can use  
 26038 any of the functions allowed for **O\_RDONLY** and **O\_WRONLY**. A message  
 26039 queue may be open multiple times in the same or different processes for  
 26040 sending messages.

26041 Any combination of the remaining flags may be specified in the value of *oflag*:

26042 **O\_CREAT** This option is used to create a message queue, and it requires two additional  
 26043 arguments: *mode*, which is of type **mode\_t**, and *attr*, which is a pointer to a  
 26044 **mq\_attr** structure. If the path name *name* has already been used to create a  
 26045 message queue that still exists, then this flag has no effect, except as noted  
 26046 under **O\_EXCL**. Otherwise, a message queue is created without any messages  
 26047 in it. The user ID of the message queue is set to the effective user ID of the  
 26048 process, and the group ID of the message queue is set to the effective group ID

|       |                     |                                                                                                          |
|-------|---------------------|----------------------------------------------------------------------------------------------------------|
| 26049 |                     | of the process. The file permission bits are set to the value of <i>mode</i> . When bits                 |
| 26050 |                     | in <i>mode</i> other than file permission bits are set, the effect is implementation-                    |
| 26051 |                     | defined. If <i>attr</i> is NULL, the message queue is created with implementation-                       |
| 26052 |                     | defined default message queue attributes. If <i>attr</i> is non-NULL and the calling                     |
| 26053 |                     | process has the appropriate privilege on <i>name</i> , the message queue <i>mq_maxmsg</i>                |
| 26054 |                     | and <i>mq_msgsize</i> attributes are set to the values of the corresponding members                      |
| 26055 |                     | in the <b>mq_attr</b> structure referred to by <i>attr</i> . If <i>attr</i> is non-NULL, but the calling |
| 26056 |                     | process does not have the appropriate privilege on <i>name</i> , the <i>mq_open()</i>                    |
| 26057 |                     | function shall fail and return an error without creating the message queue.                              |
| 26058 | O_EXCL              | If O_EXCL and O_CREAT are set, <i>mq_open()</i> fails if the message queue <i>name</i>                   |
| 26059 |                     | exists. The check for the existence of the message queue and the creation of                             |
| 26060 |                     | the message queue if it does not exist shall be atomic with respect to other                             |
| 26061 |                     | threads executing <i>mq_open()</i> naming the same <i>name</i> with O_EXCL and                           |
| 26062 |                     | O_CREAT set. If O_EXCL is set and O_CREAT is not set, the result is                                      |
| 26063 |                     | undefined.                                                                                               |
| 26064 | O_NONBLOCK          | The setting of this flag is associated with the open message queue description                           |
| 26065 |                     | and determines whether a <i>mq_send()</i> or <i>mq_receive()</i> waits for resources or                  |
| 26066 |                     | messages that are not currently available, or fails with <i>errno</i> set to [EAGAIN];                   |
| 26067 |                     | see <i>mq_send()</i> and <i>mq_receive()</i> for details.                                                |
| 26068 |                     | The <i>mq_open()</i> function does not add or remove messages from the queue.                            |
| 26069 | <b>RETURN VALUE</b> |                                                                                                          |
| 26070 |                     | Upon successful completion, the function shall return a message queue descriptor; otherwise,             |
| 26071 |                     | the function shall return ( <b>mqd_t</b> )−1 and set <i>errno</i> to indicate the error.                 |
| 26072 | <b>ERRORS</b>       |                                                                                                          |
| 26073 |                     | The <i>mq_open()</i> function shall fail if:                                                             |
| 26074 | [EACCES]            | The message queue exists and the permissions specified by <i>oflag</i> are denied, or                    |
| 26075 |                     | the message queue does not exist and permission to create the message queue                              |
| 26076 |                     | is denied.                                                                                               |
| 26077 | [EEXIST]            | O_CREAT and O_EXCL are set and the named message queue already exists.                                   |
| 26078 | [EINTR]             | The <i>mq_open()</i> function was interrupted by a signal.                                               |
| 26079 | [EINVAL]            | The <i>mq_open()</i> function is not supported for the given name.                                       |
| 26080 | [EINVAL]            | O_CREAT was specified in <i>oflag</i> , the value of <i>attr</i> is not NULL, and either                 |
| 26081 |                     | <i>mq_maxmsg</i> or <i>mq_msgsize</i> was less than or equal to zero.                                    |
| 26082 | [EMFILE]            | Too many message queue descriptors or file descriptors are currently in use by                           |
| 26083 |                     | this process.                                                                                            |
| 26084 | [ENAMETOOLONG]      |                                                                                                          |
| 26085 |                     | The length of the <i>name</i> argument exceeds {PATH_MAX} or a path name                                 |
| 26086 |                     | component is longer than {NAME_MAX}.                                                                     |
| 26087 | [ENFILE]            | Too many message queues are currently open in the system.                                                |
| 26088 | [ENOENT]            | O_CREAT is not set and the named message queue does not exist.                                           |
| 26089 | [ENOSPC]            | There is insufficient space for the creation of the new message queue.                                   |

26090 **EXAMPLES**

26091 None.

26092 **APPLICATION USAGE**

26093 None.

26094 **RATIONALE**

26095 None.

26096 **FUTURE DIRECTIONS**

26097 None.

26098 **SEE ALSO**

26099 *mq\_close()*, *mq\_getattr()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*, *mq\_timedreceive()*, *mq\_timedsend()*,  
26100 *mq\_unlink()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of  
26101 IEEE Std. 1003.1-200x, <mqqueue.h>

26102 **CHANGE HISTORY**

26103 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

26104 **Issue 6**26105 The *mq\_open()* function is marked as part of the Message Passing option.

26106 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
26107 implementation does not support the Message Passing option.

26108 The *mq\_timedreceive()* and *mq\_timedsend()* functions are added to the SEE ALSO section for  
26109 alignment with IEEE Std. 1003.1d-1999.

26110 The DESCRIPTION of O\_EXCL is updated in response to IEEE PASC Interpretation 1003.1c #48.

## 26111 NAME

26112 mq\_receive, mq\_timedreceive — receive a message from a message queue (**REALTIME**)

## 26113 SYNOPSIS

26114 MSG #include <mqueue.h>

```
26115 ssize_t mq_receive(mqd_t mqdes, char *msg_ptr, size_t msg_len,
26116 unsigned *msg_prio);
```

26117

26118 MSG TMO #include <mqueue.h>

26119 #include <time.h>

```
26120 ssize_t mq_timedreceive(mqd_t mqdes, char *restrict msg_ptr,
26121 size_t msg_len, unsigned *restrict msg_prio,
26122 const struct timespec *restrict abs_timeout);
```

26123

## 26124 DESCRIPTION

26125 The *mq\_receive()* function is used to receive the oldest of the highest priority message(s) from the  
 26126 message queue specified by *mqdes*. If the size of the buffer in bytes, specified by the *msg\_len*  
 26127 argument, is less than the *mq\_msgsize* attribute of the message queue, the function shall fail and  
 26128 return an error. Otherwise, the selected message is removed from the queue and copied to the  
 26129 buffer pointed to by the *msg\_ptr* argument.

26130 If the value of *msg\_len* is greater than {SSIZE\_MAX}, the result is implementation-defined.

26131 If the argument *msg\_prio* is not NULL, the priority of the selected message is stored in the  
 26132 location referenced by *msg\_prio*.

26133 If the specified message queue is empty and O\_NONBLOCK is not set in the message queue  
 26134 description associated with *mqdes*, *mq\_receive()* blocks until a message is enqueued on the  
 26135 message queue or until *mq\_receive()* is interrupted by a signal. If more than one thread is waiting  
 26136 to receive a message when a message arrives at an empty queue and the Priority Scheduling  
 26137 option is supported, then the thread of highest priority that has been waiting the longest shall be  
 26138 selected to receive the message. Otherwise, it is unspecified which waiting thread receives the  
 26139 message. If the specified message queue is empty and O\_NONBLOCK is set in the message  
 26140 queue description associated with *mqdes*, no message is removed from the queue, and  
 26141 *mq\_receive()* shall return an error.

26142 TMO The *mq\_timedreceive()* function is used to receive the oldest of the highest priority messages from  
 26143 the message queue specified by *mqdes* as in the *mq\_receive()* function. However, if  
 26144 O\_NONBLOCK was not specified when the message queue was opened via the *mq\_open()*  
 26145 function, and no message exists on the queue to satisfy the receive, the wait for such a message  
 26146 will be terminated when the specified timeout expires. If O\_NONBLOCK is set, this function  
 26147 shall behave identically to *mq\_receive()*.

26148 The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the  
 26149 clock on which timeouts are based (that is, when the value of that clock equals or exceeds  
 26150 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time  
 26151 of the call. If the Timers option is supported, the timeout is based on the CLOCK\_REALTIME  
 26152 clock; if the Timers option is not supported, the timeout is based on the system clock as returned  
 26153 by the *time()* function. The resolution of the timeout is the resolution of the clock on which it is  
 26154 based. The *timespec* argument is defined as a structure in the <time.h> header.

26155 Under no circumstance shall the operation fail with a timeout if a message can be removed from  
 26156 the message queue immediately. The validity of the *abs\_timeout* parameter need not be checked  
 26157 if a message can be removed from the message queue immediately.

26158 **RETURN VALUE**

26159 TMO Upon successful completion, the *mq\_receive()* and *mq\_timedreceive()* functions shall return the  
 26160 length of the selected message in bytes and the message shall be removed from the queue.  
 26161 Otherwise, no message shall be removed from the queue, the functions shall return a value of -1,  
 26162 and set *errno* to indicate the error.

26163 **ERRORS**

26164 TMO The *mq\_receive()* and *mq\_timedreceive()* functions shall fail if:

26165 [EAGAIN] O\_NONBLOCK was set in the message description associated with *mqdes*,  
 26166 and the specified message queue is empty.

26167 [EBADF] The *mqdes* argument is not a valid message queue descriptor open for reading.

26168 [EMSGSIZE] The specified message buffer size, *msg\_len*, is less than the message size  
 26169 attribute of the message queue.

26170 TMO [EINTR] The *mq\_receive()* or *mq\_timedreceive()* operation was interrupted by a signal.

26171 TMO [EINVAL] The process or thread would have blocked, and the *abs\_timeout* parameter  
 26172 specified a nanoseconds field value less than zero or greater than or equal to  
 26173 1 000 million.

26174 TMO [ETIMEDOUT] The O\_NONBLOCK flag was not set when the message queue was opened,  
 26175 but no message arrived on the queue before the specified timeout expired.

26176 TMO The *mq\_receive()* and *mq\_timedreceive()* functions may fail if:

26177 [EBADMSG] The implementation has detected a data corruption problem with the  
 26178 message.

26179 **EXAMPLES**

26180 None.

26181 **APPLICATION USAGE**

26182 None.

26183 **RATIONALE**

26184 None.

26185 **FUTURE DIRECTIONS**

26186 None.

26187 **SEE ALSO**

26188 *mq\_open()*, *mq\_send()*, *mq\_timedsend()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, *time()*, the Base  
 26189 Definitions volume of IEEE Std. 1003.1-200x, <**mqqueue.h**>, <**time.h**>

26190 **CHANGE HISTORY**

26191 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

26192 **Issue 6**

26193 The *mq\_receive()* function is marked as part of the Message Passing option.

26194 The Open Group corrigenda item U021/4 has been applied. The DESCRIPTION is changed to  
 26195 refer to *msg\_len* rather than *maxsize*.

26196 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 26197 implementation does not support the Message Passing option.

26198 The following new requirements on POSIX implementations derive from alignment with the  
 26199 Single UNIX Specification:

- 26200           • In this function it is possible for the return value to exceed the range of the type **ssize\_t** (since  
26201           **size\_t** has a larger range of positive values than **ssize\_t**). A sentence restricting the size of  
26202           the **size\_t** object is added to the description to resolve this conflict.
- 26203           The *mq\_timedreceive()* function is added for alignment with IEEE Std. 1003.1d-1999.
- 26204           The **restrict** keyword is added to the *mq\_timedreceive()* prototype for alignment with the  
26205           ISO/IEC 9899:1999 standard.
- 26206           IEEE PASC Interpretation 1003.1 #109 is applied, correcting the return type for *mq\_timedreceive()*  
26207           from **int** to **ssize\_t**.

## 26208 NAME

26209 mq\_send, mq\_timedsend — send a message to a message queue (**REALTIME**)

## 26210 SYNOPSIS

26211 MSG #include &lt;mqueue.h&gt;

26212 int mq\_send(mqd\_t mqdes, const char \*msg\_ptr, size\_t msg\_len,  
26213 unsigned msg\_prio);

26214

26215 MSG TMO #include &lt;mqueue.h&gt;

26216 #include &lt;time.h&gt;

26217 int mq\_timedsend(mqd\_t mqdes, const char \*msg\_ptr, size\_t msg\_len,  
26218 unsigned msg\_prio, const struct timespec \*abs\_timeout);

26219

## 26220 DESCRIPTION

26221 The *mq\_send()* function shall add the message pointed to by the argument *msg\_ptr* to the  
26222 message queue specified by *mqdes*. The *msg\_len* argument specifies the length of the message in  
26223 bytes pointed to by *msg\_ptr*. The value of *msg\_len* is less than or equal to the *mq\_msgsize*  
26224 attribute of the message queue, or *mq\_send()* shall fail.

26225 If the specified message queue is not full, *mq\_send()* behaves as if the message shall be inserted  
26226 into the message queue at the position indicated by the *msg\_prio* argument. A message with a  
26227 larger numeric value of *msg\_prio* shall be inserted before messages with lower values of  
26228 *msg\_prio*. A message shall be inserted after other messages in the queue, if any, with equal  
26229 *msg\_prio*. The value of *msg\_prio* shall be less than {MQ\_PRIO\_MAX}.

26230 If the specified message queue is full and O\_NONBLOCK is not set in the message queue  
26231 description associated with *mqdes*, *mq\_send()* shall block until space becomes available to  
26232 enqueue the message, or until *mq\_send()* is interrupted by a signal. If more than one thread is  
26233 waiting to send when space becomes available in the message queue and the Priority Scheduling  
26234 option is supported, then the thread of the highest priority that has been waiting the longest  
26235 shall be unblocked to send its message. Otherwise, it is unspecified which waiting thread is  
26236 unblocked. If the specified message queue is full and O\_NONBLOCK is set in the message  
26237 queue description associated with *mqdes*, the message shall not be queued and *mq\_send()* shall  
26238 return an error.

26239 TMO The *mq\_timedsend()* function adds a message to the message queue specified by *mqdes* in the  
26240 manner defined for the *mq\_send()* function. However, if the specified message queue is full and  
26241 O\_NONBLOCK is not set in the message queue description associated with *mqdes*, the wait for  
26242 sufficient room in the queue shall be terminated when the specified timeout expires. If  
26243 O\_NONBLOCK is set in the message queue description, this function shall behave identically to  
26244 *mq\_send()*.

26245 The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the  
26246 clock on which timeouts are based (that is, when the value of that clock equals or exceeds  
26247 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time  
26248 of the call. If the Timers option is supported, the timeout is based on the CLOCK\_REALTIME  
26249 clock; if the Timers option is not supported, the timeout is based on the system clock as returned  
26250 by the *time()* function. The resolution of the timeout is the resolution of the clock on which it is  
26251 based. The *timespec* argument is defined as a structure in the <time.h> header.

26252 Under no circumstance shall the operation fail with a timeout if there is sufficient room in the  
26253 queue to add the message immediately. The validity of the *abs\_timeout* parameter need not be  
26254 checked when there is sufficient room in the queue.

26255 **RETURN VALUE**

26256 TMO Upon successful completion, the `mq_send()` and `mq_timedsend()` functions shall return a value of  
 26257 zero. Otherwise, no message shall be enqueued, the functions shall return `-1`, and `errno` shall be  
 26258 set to indicate the error.

26259 **ERRORS**

26260 TMO The `mq_send()` and `mq_timedsend()` functions shall fail if:

26261 [EAGAIN] The `O_NONBLOCK` flag is set in the message queue description associated  
 26262 with `mqdes`, and the specified message queue is full.

26263 [EBADF] The `mqdes` argument is not a valid message queue descriptor open for writing.

26264 TMO [EINTR] A signal interrupted the call to `mq_send()` or `mq_timedsend()`.

26265 [EINVAL] The value of `msg_prio` was outside the valid range.

26266 TMO [EINVAL] The process or thread would have blocked, and the `abs_timeout` parameter  
 26267 specified a nanoseconds field value less than zero or greater than or equal to  
 26268 1 000 million.

26269 [EMSGSIZE] The specified message length, `msg_len`, exceeds the message size attribute of  
 26270 the message queue.

26271 TMO [ETIMEDOUT] The `O_NONBLOCK` flag was not set when the message queue was opened,  
 26272 but the timeout expired before the message could be added to the queue.

26273 **EXAMPLES**

26274 None.

26275 **APPLICATION USAGE**

26276 The value of the symbol `{MQ_PRIO_MAX}` limits the number of priority levels supported by the  
 26277 application. Message priorities range from 0 to `{MQ_PRIO_MAX}-1`.

26278 **RATIONALE**

26279 None.

26280 **FUTURE DIRECTIONS**

26281 None.

26282 **SEE ALSO**

26283 `mq_open()`, `mq_receive()`, `mq_setattr()`, `mq_timedreceive()`, `time()`, the Base Definitions volume of  
 26284 IEEE Std. 1003.1-200x, `<mqqueue.h>`, `<time.h>`

26285 **CHANGE HISTORY**

26286 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

26287 **Issue 6**

26288 The `mq_send()` function is marked as part of the Message Passing option.

26289 The `[ENOSYS]` error condition has been removed as stubs need not be provided if an  
 26290 implementation does not support the Message Passing option.

26291 The `mq_timedsend()` function is added for alignment with IEEE Std. 1003.1d-1999.

26292 **NAME**26293 mq\_setattr — set message queue attributes (**REALTIME**)26294 **SYNOPSIS**

26295 MSG #include &lt;mqqueue.h&gt;

```
26296 int mq_setattr(mqd_t mqdes, const struct mq_attr *restrict mqstat,
26297 struct mq_attr *restrict omqstat);
26298
```

26299 **DESCRIPTION**26300 The *mq\_setattr()* function is used to set attributes associated with the open message queue description referenced by the message queue descriptor specified by *mqdes*.26302 The message queue attributes corresponding to the following members defined in the **mq\_attr** structure are set to the specified values upon successful completion of *mq\_setattr()*:26304 *mq\_flags* The value of this member is the bitwise-logical OR of zero or more of  
26305 O\_NONBLOCK and any implementation-defined flags.26306 The values of the *mq\_maxmsg*, *mq\_msgsize*, and *mq\_curmsgs* members of the **mq\_attr** structure are  
26307 ignored by *mq\_setattr()*.26308 If *omqstat* is non-NULL, the *mq\_setattr()* function stores, in the location referenced by *omqstat*, the  
26309 previous message queue attributes and the current queue status. These values are the same as  
26310 would be returned by a call to *mq\_getattr()* at that point.26311 **RETURN VALUE**26312 Upon successful completion, the function shall return a value of zero and the attributes of the  
26313 message queue shall have been changed as specified.26314 Otherwise, the message queue attributes shall be unchanged, and the function shall return a  
26315 value of -1 and set *errno* to indicate the error.26316 **ERRORS**26317 The *mq\_setattr()* function shall fail if:26318 [EBADF] The *mqdes* argument is not a valid message queue descriptor.26319 **EXAMPLES**

26320 None.

26321 **APPLICATION USAGE**

26322 None.

26323 **RATIONALE**

26324 None.

26325 **FUTURE DIRECTIONS**

26326 None.

26327 **SEE ALSO**26328 *mq\_open()*, *mq\_send()*, *mq\_timedsend()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base  
26329 Definitions volume of IEEE Std. 1003.1-200x, <**mqqueue.h**>26330 **CHANGE HISTORY**

26331 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

26332 **Issue 6**

26333 The *mq\_setattr()* function is marked as part of the Message Passing option.

26334 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
26335 implementation does not support the Message Passing option.

26336 The *mq\_timedsend()* function is added to the SEE ALSO section for alignment with  
26337 IEEE Std. 1003.1d-1999.

26338 The **restrict** keyword is added to the *mq\_setattr()* prototype for alignment with the  
26339 ISO/IEC 9899:1999 standard.

26340 **NAME**26341 mq\_timedreceive — receive a message from a message queue (**REALTIME**)26342 **SYNOPSIS**

26343 MSG TMO #include &lt;mqueue.h&gt;

26344 #include &lt;time.h&gt;

26345 int mq\_timedreceive(mqd\_t mqdes, char \*restrict msg\_ptr,

26346 size\_t msg\_len, unsigned \*restrict msg\_prio,

26347 const struct timespec \*restrict abs\_timeout);

26348

26349 **DESCRIPTION**26350 Refer to *mq\_receive()*.

26351 **NAME**

26352 mq\_timedsend — send a message to a message queue (**REALTIME**)

26353 **SYNOPSIS**

26354 MSG TMO #include <mqueue.h>

26355 #include <time.h>

26356 int mq\_timedsend(mqd\_t mqdes, const char \*msg\_ptr, size\_t msg\_len,

26357 unsigned msg\_prio, const struct timespec \*abs\_timeout);

26358

26359 **DESCRIPTION**

26360 Refer to *mq\_send()*.

26361 **NAME**26362 mq\_unlink — remove a message queue (**REALTIME**)26363 **SYNOPSIS**

26364 MSG #include &lt;mqueue.h&gt;

26365 int mq\_unlink(const char \*name);

26366

26367 **DESCRIPTION**

26368 The *mq\_unlink()* function shall remove the message queue named by the path name *name*. After  
 26369 a successful call to *mq\_unlink()* with *name*, a call to *mq\_open()* with *name* fails if the flag  
 26370 O\_CREAT is not set in *flags*. If one or more processes have the message queue open when  
 26371 *mq\_unlink()* is called, destruction of the message queue is postponed until all references to the  
 26372 message queue have been closed.

26373 Calls to *mq\_open()* to recreate the message queue may fail until the message queue is actually  
 26374 removed. However, the *mq\_unlink()* call need not block until all references have been closed; it  
 26375 may return immediately.

26376 **RETURN VALUE**

26377 Upon successful completion, the function shall return a value of zero. Otherwise, the named  
 26378 message queue shall be unchanged by this function call, and the function shall return a value of  
 26379 -1 and set *errno* to indicate the error.

26380 **ERRORS**26381 The *mq\_unlink()* function shall fail if:

26382 [EACCES] Permission is denied to unlink the named message queue.

26383 [ENAMETOOLONG]

26384 The length of the *name* argument exceeds {PATH\_MAX} or a path name  
 26385 component is longer than {NAME\_MAX}.

26386 [ENOENT] The named message queue does not exist.

26387 **EXAMPLES**

26388 None.

26389 **APPLICATION USAGE**

26390 None.

26391 **RATIONALE**

26392 None.

26393 **FUTURE DIRECTIONS**

26394 None.

26395 **SEE ALSO**

26396 *mq\_close()*, *mq\_open()*, *msgctl()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of  
 26397 IEEE Std. 1003.1-200x, <mqueue.h>

26398 **CHANGE HISTORY**

26399 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

26400 **Issue 6**26401 The *mq\_unlink()* function is marked as part of the Message Passing option.

26402 The Open Group corrigenda item U021/5 has been applied, clarifying that upon unsuccessful  
 26403 completion, the named message queue is unchanged by this function.

26404  
26405

The [ENOSYS] error condition has been removed as stubs need not be provided if an implementation does not support the Message Passing option.

26406 **NAME**

26407       mrand48 — generate uniformly distributed pseudo-random signed long integers

26408 **SYNOPSIS**

26409 xSI     #include <stdlib.h>

26410       long mrand48(void);

26411

26412 **DESCRIPTION**

26413       Refer to *drand48()*.

## 26414 NAME

26415 msgctl — XSI message control operations

## 26416 SYNOPSIS

26417 XSI #include &lt;sys/msg.h&gt;

26418 int msgctl(int *msqid*, int *cmd*, struct *msqid\_ds* \**buf*);

26419

## 26420 DESCRIPTION

26421 The *msgctl()* function operates on XSI message queues (see the Base Definitions volume of  
 26422 IEEE Std. 1003.1-200x, Section 3.226, Message Queue). It is unspecified whether this function  
 26423 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 26424 page 543).

26425 The *msgctl()* function shall provide message control operations as specified by *cmd*. The  
 26426 following values for *cmd*, and the message control operations they specify, are:

26427 IPC\_STAT Place the current value of each member of the **msqid\_ds** data structure  
 26428 associated with *msqid* into the structure pointed to by *buf*. The contents of this  
 26429 structure are defined in <sys/msg.h>.

26430 IPC\_SET Set the value of the following members of the **msqid\_ds** data structure  
 26431 associated with *msqid* to the corresponding value found in the structure  
 26432 pointed to by *buf*:

26433 msg\_perm.uid  
 26434 msg\_perm.gid  
 26435 msg\_perm.mode  
 26436 msg\_qbytes

26437 IPC\_SET can only be executed by a process with appropriate privileges or that  
 26438 has an effective user ID equal to the value of **msg\_perm.cuid** or  
 26439 **msg\_perm.uid** in the **msqid\_ds** data structure associated with *msqid*. Only a  
 26440 process with appropriate privileges can raise the value of *msg\_qbytes*.

26441 IPC\_RMID Remove the message queue identifier specified by *msqid* from the system and  
 26442 destroy the message queue and **msqid\_ds** data structure associated with it.  
 26443 IPC\_RMID can only be executed by a process with appropriate privileges or  
 26444 one that has an effective user ID equal to the value of **msg\_perm.cuid** or  
 26445 **msg\_perm.uid** in the **msqid\_ds** data structure associated with *msqid*.

## 26446 RETURN VALUE

26447 Upon successful completion, *msgctl()* shall return 0; otherwise, it shall return -1 and set *errno* to  
 26448 indicate the error.

## 26449 ERRORS

26450 The *msgctl()* function shall fail if:

26451 [EACCES] The argument *cmd* is IPC\_STAT and the calling process does not have read  
 26452 permission; see Section 2.7 (on page 541).

26453 [EINVAL] The value of *msqid* is not a valid message queue identifier; or the value of *cmd*  
 26454 is not a valid command.

26455 [EPERM] The argument *cmd* is IPC\_RMID or IPC\_SET and the effective user ID of the  
 26456 calling process is not equal to that of a process with appropriate privileges  
 26457 and it is not equal to the value of **msg\_perm.cuid** or **msg\_perm.uid** in the data  
 26458 structure associated with *msqid*.

26459 [EPERM] The argument *cmd* is `IPC_SET`, an attempt is being made to increase to the  
26460 value of *msg\_qbytes*, and the effective user ID of the calling process does not  
26461 have appropriate privileges.

26462 **EXAMPLES**

26463 None.

26464 **APPLICATION USAGE**

26465 The POSIX Realtime Extension defines alternative interfaces for interprocess communication  
26466 (IPC). Application developers who need to use IPC should design their applications so that  
26467 modules using the IPC routines described in Section 2.7 (on page 541) can be easily modified to  
26468 use the alternative interfaces.

26469 **RATIONALE**

26470 None.

26471 **FUTURE DIRECTIONS**

26472 None.

26473 **SEE ALSO**

26474 *mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*,  
26475 *mq\_unlink()*, *msgget()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
26476 `<sys/msg.h>`, Section 2.7 (on page 541)

26477 **CHANGE HISTORY**

26478 First released in Issue 2. Derived from Issue 2 of the SVID.

26479 **Issue 4**

26480 The function is no longer marked as OPTIONAL FUNCTIONALITY.

26481 Inclusion of the `<sys/types.h>` and `<sys/ipc.h>` headers is removed from the SYNOPSIS section.

26482 A FUTURE DIRECTIONS section is added warning application developers about migration to  
26483 IEEE 1003.4 interfaces for interprocess communication.

26484 The [ENOSYS] error is removed from the ERRORS section.

26485 **Issue 5**

26486 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
26487 DIRECTIONS to a new APPLICATION USAGE section.

## 26488 NAME

26489 msgget — get the XSI message queue identifier

## 26490 SYNOPSIS

26491 XSI #include &lt;sys/msg.h&gt;

26492 int msgget(key\_t key, int msgflg);

26493

## 26494 DESCRIPTION

26495 The *msgget()* function operates on XSI message queues (see the Base Definitions volume of  
 26496 IEEE Std. 1003.1-200x, Section 3.226, Message Queue). It is unspecified whether this function  
 26497 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 26498 page 543).

26499 The *msgget()* function shall return the message queue identifier associated with the argument  
 26500 *key*.

26501 A message queue identifier, associated message queue, and data structure (see <sys/msg.h>), are  
 26502 created for the argument *key* if one of the following is true:

- 26503 • The argument *key* is equal to `IPC_PRIVATE`.
- 26504 • The argument *key* does not already have a message queue identifier associated with it, and  
 26505 (*msgflg* & `IPC_CREAT`) is non-zero.

26506 Upon creation, the data structure associated with the new message queue identifier is initialized  
 26507 as follows:

- 26508 • `msg_perm.cuid`, `msg_perm.uid`, `msg_perm.cgid`, and `msg_perm.gid` are set equal to the  
 26509 effective user ID and effective group ID, respectively, of the calling process.
- 26510 • The low-order 9 bits of `msg_perm.mode` are set equal to the low-order 9 bits of *msgflg*.
- 26511 • `msg_qnum`, `msg_lspid`, `msg_lrpid`, `msg_stime`, and `msg_rtime` are set equal to 0.
- 26512 • `msg_ctime` is set equal to the current time.
- 26513 • `msg_qbytes` is set equal to the system limit.

## 26514 RETURN VALUE

26515 Upon successful completion, *msgget()* shall return a non-negative integer, namely a message  
 26516 queue identifier. Otherwise, it shall return `-1` and set *errno* to indicate the error.

## 26517 ERRORS

26518 The *msgget()* function shall fail if:

- |                         |          |                                                                                                                                                                                                           |
|-------------------------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 26519<br>26520<br>26521 | [EACCES] | A message queue identifier exists for the argument <i>key</i> , but operation<br>permission as specified by the low-order 9 bits of <i>msgflg</i> would not be granted;<br>see Section 2.7 (on page 541). |
| 26522<br>26523          | [EEXIST] | A message queue identifier exists for the argument <i>key</i> but (( <i>msgflg</i> &<br><code>IPC_CREAT</code> ) && ( <i>msgflg</i> & <code>IPC_EXCL</code> )) is non-zero.                               |
| 26524<br>26525          | [ENOENT] | A message queue identifier does not exist for the argument <i>key</i> and ( <i>msgflg</i> &<br><code>IPC_CREAT</code> ) is 0.                                                                             |
| 26526<br>26527<br>26528 | [ENOSPC] | A message queue identifier is to be created but the system-imposed limit on<br>the maximum number of allowed message queue identifiers system-wide<br>would be exceeded.                                  |

**26529 EXAMPLES**

26530 None.

**26531 APPLICATION USAGE**

26532 The POSIX Realtime Extension defines alternative interfaces for interprocess communication  
26533 (IPC). Application developers who need to use IPC should design their applications so that  
26534 modules using the IPC routines described in Section 2.7 (on page 541) can be easily modified to  
26535 use the alternative interfaces.

**26536 RATIONALE**

26537 None.

**26538 FUTURE DIRECTIONS**

26539 None.

**26540 SEE ALSO**

26541 *mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*,  
26542 *mq\_unlink()*, *msgctl()*, *msgrcv()*, *msgsnd()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
26543 `<sys/msg.h>`, Section 2.7 (on page 541)

**26544 CHANGE HISTORY**

26545 First released in Issue 2. Derived from Issue 2 of the SVID.

**26546 Issue 4**

26547 The function is no longer marked as OPTIONAL FUNCTIONALITY.

26548 Inclusion of the `<sys/types.h>` and `<sys/ipc.h>` headers is removed from the SYNOPSIS section.

26549 The [ENOSYS] error is removed from the ERRORS section.

**26550 Issue 5**

26551 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
26552 DIRECTIONS to a new APPLICATION USAGE section.

## 26553 NAME

26554 msgrcv — XSI message receive operation

## 26555 SYNOPSIS

26556 XSI 

```
#include <sys/msg.h>
```

26557 

```
ssize_t msgrcv(int msqid, void *msgp, size_t msgsz, long msgtyp,
26558 int msgflg);
```

26559

## 26560 DESCRIPTION

26561 The *msgrcv()* function operates on XSI message queues (see the Base Definitions volume of  
26562 IEEE Std. 1003.1-200x, Section 3.226, Message Queue). It is unspecified whether this function  
26563 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
26564 page 543).

26565 The *msgrcv()* function shall read a message from the queue associated with the message queue  
26566 identifier specified by *msqid* and place it in the user-defined buffer pointed to by *msgp*.

26567 The application shall ensure that the argument *msgp* points to a user-defined buffer that contains  
26568 first a field of type **long** specifying the type of the message, and then a data portion that holds  
26569 the data bytes of the message. The structure below is an example of what this user-defined  
26570 buffer might look like:

```
26571 struct mymsg {
26572 long mtype; /* Message type. */
26573 char mtext[1]; /* Message text. */
26574 }
```

26575 The structure member *mtype* is the received message's type as specified by the sending process.

26576 The structure member *mtext* is the text of the message.

26577 The argument *msgsz* specifies the size in bytes of *mtext*. The received message is truncated to  
26578 *msgsz* bytes if it is larger than *msgsz* and (*msgflg* & MSG\_NOERROR) is non-zero. The truncated  
26579 part of the message is lost and no indication of the truncation is given to the calling process.

26580 If the value of *msgsz* is greater than {SSIZE\_MAX}, the result is implementation-defined.

26581 The argument *msgtyp* specifies the type of message requested as follows:

- 26582 • If *msgtyp* is 0, the first message on the queue is received.
- 26583 • If *msgtyp* is greater than 0, the first message of type *msgtyp* is received.
- 26584 • If *msgtyp* is less than 0, the first message of the lowest type that is less than or equal to the  
26585 absolute value of *msgtyp* is received.

26586 The argument *msgflg* specifies the action to be taken if a message of the desired type is not on the  
26587 queue. These are as follows:

- 26588 • If (*msgflg* & IPC\_NOWAIT) is non-zero, the calling thread shall return immediately with a  
26589 return value of -1 and *errno* set to [ENOMSG].
- 26590 • If (*msgflg* & IPC\_NOWAIT) is 0, the calling thread shall suspend execution until one of the  
26591 following occurs:
  - 26592 — A message of the desired type is placed on the queue.
  - 26593 — The message queue identifier *msqid* is removed from the system; when this occurs, *errno*  
26594 shall be set equal to [EIDRM] and -1 shall be returned.

26595 — The calling thread receives a signal that is to be caught; in this case a message is not  
 26596 received and the calling thread resumes execution in the manner prescribed in *sigaction()*.

26597 Upon successful completion, the following actions are taken with respect to the data structure  
 26598 associated with *msqid*:

- 26599 • **msg\_qnum** is decremented by 1.
- 26600 • **msg\_lrpri** is set equal to the process ID of the calling process.
- 26601 • **msg\_rtime** is set equal to the current time.

#### 26602 RETURN VALUE

26603 Upon successful completion, *msgrcv()* shall return a value equal to the number of bytes actually  
 26604 placed into the buffer *mtext*. Otherwise, no message shall be received, *msgrcv()* shall return  
 26605 (**ssize\_t**)−1, and *errno* shall be set to indicate the error.

#### 26606 ERRORS

26607 The *msgrcv()* function shall fail if:

|       |          |                                                                                                 |  |
|-------|----------|-------------------------------------------------------------------------------------------------|--|
| 26608 | [E2BIG]  | The value of <i>mtext</i> is greater than <i>msgsz</i> and ( <i>msgflg</i> & MSG_NOERROR) is 0. |  |
| 26609 | [EACCES] | Operation permission is denied to the calling process; see Section 2.7 (on page                 |  |
| 26610 |          | 541).                                                                                           |  |
| 26611 | [EIDRM]  | The message queue identifier <i>msqid</i> is removed from the system.                           |  |
| 26612 | [EINTR]  | The <i>msgrcv()</i> function was interrupted by a signal.                                       |  |
| 26613 | [EINVAL] | <i>msqid</i> is not a valid message queue identifier.                                           |  |
| 26614 | [ENOMSG] | The queue does not contain a message of the desired type and ( <i>msgflg</i> &                  |  |
| 26615 |          | IPC_NOWAIT) is non-zero.                                                                        |  |

#### 26616 EXAMPLES

##### 26617 Receiving a Message

26618 The following example receives the first message on the queue (based on the value of the  
 26619 *msgtype* argument, 0). The queue is identified by the *msqid* argument (assuming that the value  
 26620 has previously been set). This call specifies that an error should be reported if no message is  
 26621 available, but not if the message is too large. The message size is calculated directly using the  
 26622 *sizeof* operator.

```

26623 #include <sys/msg.h>
26624 ...
26625 int result;
26626 int msqid;
26627 struct message {
26628 long type;
26629 char text[20];
26630 } msg;
26631 long msgtyp = 0;
26632 ...
26633 result = msgrcv(msqid, (void *) &msg, sizeof(msg.text),
26634 msgtyp, MSG_NOERROR | IPC_NOWAIT);

```

**26635 APPLICATION USAGE**

26636 The POSIX Realtime Extension defines alternative interfaces for interprocess communication  
26637 (IPC). Application developers who need to use IPC should design their applications so that  
26638 modules using the IPC routines described in Section 2.7 (on page 541) can be easily modified to  
26639 use the alternative interfaces.

**26640 RATIONALE**

26641 None.

**26642 FUTURE DIRECTIONS**

26643 None.

**26644 SEE ALSO**

26645 *mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*,  
26646 *mq\_unlink()*, *msgctl()*, *msgget()*, *msgsnd()*, *sigaction()*, the Base Definitions volume of  
26647 IEEE Std. 1003.1-200x, <sys/msg.h>, Section 2.7 (on page 541)

**26648 CHANGE HISTORY**

26649 First released in Issue 2. Derived from Issue 2 of the SVID.

**26650 Issue 4**

26651 The function is no longer marked as OPTIONAL FUNCTIONALITY.

26652 Inclusion of the <sys/types.h> and <sys/ipc.h> headers is removed from the SYNOPSIS section.

26653 The [ENOSYS] error is removed from the ERRORS section.

26654 A FUTURE DIRECTIONS section is added warning application developers about migration to  
26655 IEEE 1003.4 interfaces for interprocess communication.

**26656 Issue 5**

26657 The type of the return value is changed from **int** to **ssize\_t**, and a warning is added to the  
26658 DESCRIPTION about values of *msgsz* larger the {SSIZE\_MAX}.

26659 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
26660 DIRECTIONS to the APPLICATION USAGE section.

**26661 Issue 6**

26662 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

## 26663 NAME

26664 msgsnd — XSI message send operation

## 26665 SYNOPSIS

26666 XSI 

```
#include <sys/msg.h>
```

26667 

```
int msgsnd(int msqid, const void *msgp, size_t msgsz, int msgflg);
```

26668

## 26669 DESCRIPTION

26670 The *msgsnd()* function operates on XSI message queues (see the Base Definitions volume of  
 26671 IEEE Std. 1003.1-200x, Section 3.226, Message Queue). It is unspecified whether this function  
 26672 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 26673 page 543).

26674 The *msgsnd()* function is used to send a message to the queue associated with the message  
 26675 queue identifier specified by *msqid*.

26676 The application shall ensure that the argument *msgp* points to a user-defined buffer that contains  
 26677 first a field of type **long** specifying the type of the message, and then a data portion that holds  
 26678 the data bytes of the message. The structure below is an example of what this user-defined  
 26679 buffer might look like:

```
26680 struct mymsg {
26681 long mtype; /* Message type. */
26682 char mtext[1]; /* Message text. */
26683 }
```

26684 The structure member *mtype* is a non-zero positive type **long** that can be used by the receiving  
 26685 process for message selection.

26686 The structure member *mtext* is any text of length *msgsz* bytes. The argument *msgsz* can range  
 26687 from 0 to a system-imposed maximum.

26688 The argument *msgflg* specifies the action to be taken if one or more of the following are true:

- 26689 • The number of bytes already on the queue is equal to *msg\_qbytes*; see `<sys/msg.h>`.
- 26690 • The total number of messages on all queues system-wide is equal to the system-imposed  
 26691 limit.

26692 These actions are as follows:

- 26693 • If (*msgflg* & IPC\_NOWAIT) is non-zero, the message shall not be sent and the calling thread  
 26694 shall return immediately.
- 26695 • If (*msgflg* & IPC\_NOWAIT) is 0, the calling thread shall suspend execution until one of the  
 26696 following occurs:
  - 26697 — The condition responsible for the suspension no longer exists, in which case the message  
 26698 is sent.
  - 26699 — The message queue identifier *msqid* is removed from the system; when this occurs, *errno*  
 26700 shall be set equal to [EIDRM] and `-1` shall be returned.
  - 26701 — The calling thread receives a signal that is to be caught; in this case the message is not  
 26702 sent and the calling thread resumes execution in the manner prescribed in *sigaction()*.

26703 Upon successful completion, the following actions are taken with respect to the data structure  
 26704 associated with *msqid*; see `<sys/msg.h>`:

- 26705 • **msg\_qnum** is incremented by 1.
- 26706 • **msg\_lspid** is set equal to the process ID of the calling process.
- 26707 • **msg\_stime** is set equal to the current time.

**26708 RETURN VALUE**

26709 Upon successful completion, *msgsnd()* shall return 0; otherwise, no message shall be sent,  
 26710 *msgsnd()* shall return -1, and *errno* shall be set to indicate the error.

**26711 ERRORS**

26712 The *msgsnd()* function shall fail if:

- 26713 [EACCES] Operation permission is denied to the calling process; see Section 2.7 (on page  
 26714 541).
- 26715 [EAGAIN] The message cannot be sent for one of the reasons cited above and (*msgflg* &  
 26716 IPC\_NOWAIT) is non-zero.
- 26717 [EIDRM] The message queue identifier *msqid* is removed from the system.
- 26718 [EINTR] The *msgsnd()* function was interrupted by a signal.
- 26719 [EINVAL] The value of *msqid* is not a valid message queue identifier, or the value of  
 26720 *mtype* is less than 1; or the value of *msgsz* is less than 0 or greater than the  
 26721 system-imposed limit.

**26722 EXAMPLES****26723 Sending a Message**

26724 The following example sends a message to the queue identified by the *msqid* argument  
 26725 (assuming that value has previously been set). This call specifies that an error should be  
 26726 reported if no message is available. The message size is calculated directly using the *sizeof*  
 26727 operator.

```

26728 #include <sys/msg.h>
26729 ...
26730 int result;
26731 int msqid;
26732 struct message {
26733 long type;
26734 char text[20];
26735 } msg;

26736 msg.type = 1;
26737 strcpy(msg.text, "This is message 1");
26738 ...
26739 result = msgsnd(msqid, (void *) &msg, sizeof(msg.text), IPC_NOWAIT);

```

**26740 APPLICATION USAGE**

26741 The POSIX Realtime Extension defines alternative interfaces for interprocess communication  
 26742 (IPC). Application developers who need to use IPC should design their applications so that  
 26743 modules using the IPC routines described in Section 2.7 (on page 541) can be easily modified to  
 26744 use the alternative interfaces.

26745 **RATIONALE**

26746 None.

26747 **FUTURE DIRECTIONS**

26748 None.

26749 **SEE ALSO**

26750 *mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*,  
26751 *mq\_unlink()*, *msgctl()*, *msgget()*, *msgrcv()*, *sigaction()*, the Base Definitions volume of  
26752 IEEE Std. 1003.1-200x, <sys/msg.h>, Section 2.7 (on page 541)

26753 **CHANGE HISTORY**

26754 First released in Issue 2. Derived from Issue 2 of the SVID.

26755 **Issue 4**

26756 The function is no longer marked as OPTIONAL FUNCTIONALITY.

26757 Inclusion of the &lt;sys/types.h&gt; and &lt;sys/ipc.h&gt; headers is removed from the SYNOPSIS section.

26758 Also the type of argument *msgp* is changed from **void\*** to **const void\***.

26759 In the DESCRIPTION, the example of a message buffer is changed:

- 26760 • Explicitly to define the first member as being of type **long**
- 26761 • To define the size of the message array *mtext*

26762 The [ENOSYS] error is removed from the ERRORS section.

26763 A FUTURE DIRECTIONS section is added warning application developers about migration to  
26764 IEEE 1003.4 interfaces for interprocess communication.

26765 **Issue 5**

26766 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
26767 DIRECTIONS to a new APPLICATION USAGE section.

26768 **Issue 6**

26769 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

26770 **NAME**

26771 `msync` — synchronize memory with physical storage

26772 **SYNOPSIS**

26773 MF SIO `#include <sys/mman.h>`

26774 `int msync(void *addr, size_t len, int flags);`

26775

26776 **DESCRIPTION**

26777 The `msync()` function shall write all modified data to permanent storage locations, if any, in  
 26778 those whole pages containing any part of the address space of the process starting at address  
 26779 `addr` and continuing for `len` bytes. If no such storage exists, `msync()` need not have any effect. If  
 26780 requested, the `msync()` function then invalidates cached copies of data.

26781 The implementation shall require that `addr` be a multiple of the page size as returned by  
 26782 `sysconf()`.

26783 For mappings to files, the `msync()` function ensures that all write operations are completed as  
 26784 defined for synchronized I/O data integrity completion. It is unspecified whether the  
 26785 implementation also writes out other file attributes. When the `msync()` function is called on  
 26786 MAP\_PRIVATE mappings, any modified data shall not be written to the underlying object and  
 26787 shall not cause such data to be made visible to other processes. It is unspecified whether data in  
 26788 SHM|TYM MAP\_PRIVATE mappings has any permanent storage locations. The effect of `msync()` on a  
 26789 shared memory object or a typed memory object is unspecified. The behavior of this function is  
 26790 unspecified if the mapping was not established by a call to `mmap()`.

26791 The `flags` argument is constructed from the bitwise-inclusive OR of one or more of the following  
 26792 flags defined in the header `<sys/mman.h>`:

26793

26794

26795

26796

26797

| Symbolic Constant | Description                  |
|-------------------|------------------------------|
| MS_ASYNC          | Perform asynchronous writes. |
| MS_SYNC           | Perform synchronous writes.  |
| MS_INVALIDATE     | Invalidate cached data.      |

26798 When MS\_ASYNC is specified, `msync()` shall return immediately once all the write operations  
 26799 are initiated or queued for servicing; when MS\_SYNC is specified, `msync()` shall not return until  
 26800 all write operations are completed as defined for synchronized I/O data integrity completion.  
 26801 Either MS\_ASYNC or MS\_SYNC is specified, but not both.

26802 When MS\_INVALIDATE is specified, `msync()` invalidates all cached copies of mapped data that  
 26803 are inconsistent with the permanent storage locations such that subsequent references shall  
 26804 obtain data that was consistent with the permanent storage locations sometime between the call  
 26805 to `msync()` and the first subsequent memory reference to the data.

26806 If `msync()` causes any write to a file, the file's `st_ctime` and `st_mtime` fields shall be marked for  
 26807 update.

26808 **RETURN VALUE**

26809 Upon successful completion, `msync()` shall return 0; otherwise, it shall return -1 and set `errno` to  
 26810 indicate the error.

26811 **ERRORS**

26812 The `msync()` function shall fail if:

26813 [EBUSY] Some or all of the addresses in the range starting at `addr` and continuing for `len`  
 26814 bytes are locked, and MS\_INVALIDATE is specified.

- 26815 [EINVAL] The value of *flags* is invalid.
- 26816 [EINVAL] The value of *addr* is not a multiple of the page size, {PAGESIZE}.
- 26817 [ENOMEM] The addresses in the range starting at *addr* and continuing for *len* bytes are  
26818 outside the range allowed for the address space of a process or specify one or  
26819 more pages that are not mapped.

**26820 EXAMPLES**

26821 None.

**26822 APPLICATION USAGE**

26823 The *msync()* function is only supported if the Memory Mapped Files option and the  
26824 Synchronized Input and Output option are supported, and thus need not be available on all  
26825 implementations.

26826 The *msync()* function should be used by programs that require a memory object to be in a  
26827 known state; for example, in building transaction facilities.

26828 Normal system activity can cause pages to be written to disk. Therefore, there are no guarantees  
26829 that *msync()* is the only control over when pages are or are not written to disk.

**26830 RATIONALE**

26831 The *msync()* function is used to write out data in a mapped region to the permanent storage for  
26832 the underlying object. The call to *msync()* ensures data integrity of the file.

26833 After the data is written out, any cached data may be invalidated if the MS\_INVALIDATE flag  
26834 was specified. This is useful on systems that do not support read/write consistency.

**26835 FUTURE DIRECTIONS**

26836 None.

**26837 SEE ALSO**

26838 *mmap()*, *sysconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/mman.h>

**26839 CHANGE HISTORY**

26840 First released in Issue 4, Version 2.

**26841 Issue 5**

26842 Moved from X/OPEN UNIX extension to BASE.

26843 Aligned with *msync()* in the POSIX Realtime Extension as follows:

- 26844 • The DESCRIPTION is extensively reworded.
- 26845 • [EBUSY] and a new form of [EINVAL] are added as mandatory error conditions.

**26846 Issue 6**

26847 The *msync()* function is marked as part of the Memory Mapped Files and Synchronized Input  
26848 and Output options.

26849 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 26850 • The [EBUSY] mandatory error condition is added.

26851 The following new requirements on POSIX implementations derive from alignment with the  
26852 Single UNIX Specification:

- 26853 • The DESCRIPTION is updated to state that implementations require *addr* to be a multiple of  
26854 the page size.
- 26855 • The second [EINVAL] error condition is made mandatory.

26856  
26857

The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by adding reference to typed memory objects.

26858 **NAME**

26859           munlock — unlock a range of process address space

26860 **SYNOPSIS**

26861 MLR       #include &lt;sys/mman.h&gt;

26862           int munlock(const void \* *addr*, size\_t *len*);

26863

26864 **DESCRIPTION**26865           Refer to *mlock()*.

26866 **NAME**

26867           munlockall — unlock the address space of a process

26868 **SYNOPSIS**

26869 ML       #include <sys/mman.h>

26870           int munlockall(void);

26871

26872 **DESCRIPTION**

26873           Refer to *mlockall()*.

26874 **NAME**26875 `munmap` — unmap pages of memory26876 **SYNOPSIS**26877 MF|SHM `#include <sys/mman.h>`26878 `int munmap(void *addr, size_t len);`

26879

26880 **DESCRIPTION**

26881 The `munmap()` function shall remove any mappings for those entire pages containing any part of  
 26882 the address space of the process starting at `addr` and continuing for `len` bytes. Further references  
 26883 to these pages result in the generation of a SIGSEGV signal to the process. If there are no  
 26884 mappings in the specified address range, then `munmap()` has no effect.

26885 The implementation shall require that `addr` be a multiple of the page size {PAGESIZE}.

26886 If a mapping to be removed was private, any modifications made in this address range shall be  
 26887 discarded.

26888 ML|MLR Any memory locks (see `mlock()` and `mlockall()`) associated with this address range shall be  
 26889 removed, as if by an appropriate call to `munlock()`.

26890 TYM If a mapping removed from a typed memory object causes the corresponding address range of  
 26891 the memory pool to be inaccessible by any process in the system except through allocatable  
 26892 mappings (that is, mappings of typed memory objects opened with the  
 26893 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE flag), then that range of the memory pool shall  
 26894 become deallocated and may become available to satisfy future typed memory allocation  
 26895 requests.

26896 A mapping removed from a typed memory object opened with the  
 26897 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE flag shall not affect in any way the availability of  
 26898 that typed memory for allocation.

26899 The behavior of this function is unspecified if the mapping was not established by a call to  
 26900 `mmap()`.

26901 **RETURN VALUE**

26902 Upon successful completion, `munmap()` shall return 0; otherwise, it shall return -1 and set `errno`  
 26903 to indicate the error.

26904 **ERRORS**

26905 The `munmap()` function shall fail if:

26906 [EINVAL] Addresses in the range [`addr`,`addr+len`) are outside the valid range for the  
 26907 address space of a process.

26908 [EINVAL] The `len` argument is 0.

26909 [EINVAL] The `addr` argument is not a multiple of the page size as returned by `sysconf()`.

26910 **EXAMPLES**

26911 None.

26912 **APPLICATION USAGE**

26913 The *munmap()* function is only supported if the Memory Mapped Files option or the Shared  
26914 Memory Objects option is supported.

26915 **RATIONALE**26916 The *munmap()* function corresponds to SVR4, just as the *mmap()* function does.

26917 It is possible that an application has applied process memory locking to a region that contains  
26918 shared memory. If this has occurred, the *munmap()* call ignores those locks and, if necessary,  
26919 causes those locks to be removed.

26920 **FUTURE DIRECTIONS**

26921 None.

26922 **SEE ALSO**

26923 *mlock()*, *mlockall()*, *mmap()*, *posix\_typed\_mem\_open()*, *sysconf()*, the Base Definitions volume of  
26924 IEEE Std. 1003.1-200x, <signal.h>, <sys/mman.h>

26925 **CHANGE HISTORY**

26926 First released in Issue 4, Version 2.

26927 **Issue 5**

26928 Moved from X/OPEN UNIX extension to BASE.

26929 Aligned with *munmap()* in the POSIX Realtime Extension as follows:

- 26930 • The DESCRIPTION is extensively reworded.
- 26931 • The SIGBUS error is no longer permitted to be generated.

26932 **Issue 6**

26933 The *munmap()* function is marked as part of the Memory Mapped Files and Shared Memory  
26934 Objects option.

26935 The following new requirements on POSIX implementations derive from alignment with the  
26936 Single UNIX Specification:

- 26937 • The DESCRIPTION is updated to state that implementations require *addr* to be a multiple of  
26938 the page size.
- 26939 • The [EINVAL] error conditions are added.

26940 The following changes are made for alignment with IEEE Std. 1003.1j-2000:

- 26941 • Semantics for typed memory objects are added to the DESCRIPTION.
- 26942 • The *posix\_typed\_mem\_open()* function is added to the SEE ALSO section.

26943 **NAME**

26944 nan, nanf, nanl — return quiet NaN

26945 **SYNOPSIS**

26946 #include &lt;math.h&gt;

26947 double nan(const char \*tagp);

26948 float nanf(const char \*tagp);

26949 long double nanl(const char \*tagp);

26950 **DESCRIPTION**

26951 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 26952 conflict between the requirements described here and the ISO C standard is unintentional. This  
 26953 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

26954 The function call *nan("n-char-sequence")* shall be equivalent to:26955 strtod("NAN(*n-char-sequence*)", (char \*\*) NULL);26956 The function call *nan("")* shall be equivalent to:

26957 strtod("NAN()", (char \*\*) NULL)

26958 If *tagp* does not point to an *n-char* sequence or an empty string, the function call shall be  
 26959 equivalent to:

26960 strtod("NAN", (char \*\*) NULL)

26961 Function calls to *nanf()* and *nanl()* are equivalent to the corresponding function calls to *strtof()*  
 26962 and *strtold()*.

26963 **RETURN VALUE**26964 These functions shall return a quiet NaN, if available, with content indicated through *tagp*.

26965 If the implementation does not support quiet NaNs, these functions shall return zero.

26966 **ERRORS**

26967 No errors are defined.

26968 **EXAMPLES**

26969 None.

26970 **APPLICATION USAGE**

26971 None.

26972 **RATIONALE**

26973 None.

26974 **FUTURE DIRECTIONS**

26975 None.

26976 **SEE ALSO**

26977 *strtod()*, <REFERENCE UNDEFINED>(strtold), the Base Definitions volume of  
 26978 IEEE Std. 1003.1-200x, <math.h>

26979 **CHANGE HISTORY**

26980 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

26981 **NAME**26982 nanosleep — high resolution sleep (**REALTIME**)26983 **SYNOPSIS**26984 TMR `#include <time.h>`26985 `int nanosleep(const struct timespec *rqtp, struct timespec *rmtp);`

26986

26987 **DESCRIPTION**

26988 The *nanosleep()* function shall cause the current thread to be suspended from execution until  
 26989 either the time interval specified by the *rqtp* argument has elapsed or a signal is delivered to the  
 26990 calling thread, and its action is to invoke a signal-catching function or to terminate the process.  
 26991 The suspension time may be longer than requested because the argument value is rounded up to  
 26992 an integer multiple of the sleep resolution or because of the scheduling of other activity by the  
 26993 system. But, except for the case of being interrupted by a signal, the suspension time shall not be  
 26994 less than the time specified by *rqtp*, as measured by the system clock, **CLOCK\_REALTIME**.

26995 The use of the *nanosleep()* function has no effect on the action or blockage of any signal.26996 **RETURN VALUE**26997 If the *nanosleep()* function returns because the requested time has elapsed, its return value shall  
26998 be zero.

26999 If the *nanosleep()* function returns because it has been interrupted by a signal, the function shall  
 27000 return a value of  $-1$  and set *errno* to indicate the interruption. If the *rmtp* argument is non-NULL,  
 27001 the **timespec** structure referenced by it is updated to contain the amount of time remaining in  
 27002 the interval (the requested time minus the time actually slept). If the *rmtp* argument is NULL, the  
 27003 remaining time is not returned.

27004 If *nanosleep()* fails, it shall return a value of  $-1$  and set *errno* to indicate the error.27005 **ERRORS**27006 The *nanosleep()* function shall fail if:27007 [EINTR] The *nanosleep()* function was interrupted by a signal.27008 [EINVAL] The *rqtp* argument specified a nanosecond value less than zero or greater than  
27009 or equal to 1000 million.27010 **EXAMPLES**

27011 None.

27012 **APPLICATION USAGE**

27013 None.

27014 **RATIONALE**

27015 It is common to suspend execution of a process for an interval in order to poll the status of a  
 27016 non-interrupting function. A large number of actual needs can be met with a simple extension to  
 27017 *sleep()* that provides finer resolution.

27018 In the POSIX.1-1990 standard and SVR4, it is possible to implement such a routine, but the  
 27019 frequency of wakeup is limited by the resolution of the *alarm()* and *sleep()* functions. In 4.3 BSD,  
 27020 it is possible to write such a routine using no static storage and reserving no system facilities.  
 27021 Although it is possible to write a function with similar functionality to *sleep()* using the  
 27022 remainder of the timers function, such a function requires the use of signals and the reservation  
 27023 of some signal number. This volume of IEEE Std. 1003.1-200x requires that *nanosleep()* be non-  
 27024 intrusive of the signals function.

27025 The *nanosleep()* function shall return a value of 0 on success and –1 on failure or if interrupted.  
27026 This latter case is different from *sleep()*. This was done because the remaining time is returned  
27027 via an argument structure pointer, *rmtp*, instead of as the return value.

27028 **FUTURE DIRECTIONS**

27029 None.

27030 **SEE ALSO**

27031 *sleep()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

27032 **CHANGE HISTORY**

27033 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

27034 **Issue 6**

27035 The *nanosleep()* function is marked as part of the Timers option.

27036 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
27037 implementation does not support the Timers option.

27038 **NAME**

27039 nearbyint, nearbyintf, nearbyintl — floating-point rounding functions

27040 **SYNOPSIS**

27041 #include <math.h>

27042 double nearbyint(double x);

27043 float nearbyintf(float x);

27044 long double nearbyintl(long double x);

27045 **DESCRIPTION**

27046 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
27047 conflict between the requirements described here and the ISO C standard is unintentional. This  
27048 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

27049 These functions shall round their argument to an integer value in floating-point format, using  
27050 the current rounding direction and without raising the inexact floating-point exception.

27051 An application wishing to check for error situations should set *errno* to 0 before calling these  
27052 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

27053 **RETURN VALUE**

27054 Upon successful completion, these functions shall return the rounded integer value.

27055 If *x* is  $\pm\text{Inf}$ , these functions shall return *x*.

27056 If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

27057 **ERRORS**

27058 These functions may fail if:

27059 [EDOM] The value of *x* is NaN.

27060 **EXAMPLES**

27061 None.

27062 **APPLICATION USAGE**

27063 None.

27064 **RATIONALE**

27065 None.

27066 **FUTURE DIRECTIONS**

27067 None.

27068 **SEE ALSO**

27069 The Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

27070 **CHANGE HISTORY**

27071 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

27072 **NAME**

27073 nextafter, nextafterf, nextafterl, nexttoward, nexttowardf, nexttowardl — next representable  
 27074 double-precision floating-point number

27075 **SYNOPSIS**

```
27076 xSI #include <math.h>
27077 double nextafter(double x, double y);
27078 float nextafterf(float x, float y);
27079 long double nextafterl(long double x, long double y);
27080 double nexttoward(double x, long double y);
27081 float nexttowardf(float x, long double y);
27082 long double nexttowardl(long double x, long double y);
27083
```

27084 **DESCRIPTION**

27085 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 27086 conflict between the requirements described here and the ISO C standard is unintentional. This  
 27087 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

27088 The *nextafter()*, *nextafterf()*, and *nextafterl()* functions shall compute the next representable  
 27089 double-precision floating-point value following *x* in the direction of *y*. Thus, if *y* is less than *x*,  
 27090 *nextafter()* returns the largest representable floating-point number less than *x*. The *nextafter()*,  
 27091 *nextafterf()*, and *nextafterl()* functions shall return *y* if *x* equals *y*.

27092 The *nexttoward()*, *nexttowardf()*, and *nexttowardl()* functions are equivalent to the corresponding  
 27093 *nextafter()* functions, except that the second parameter has type **long double** and the functions  
 27094 return *y* converted to the type of the function if *x* equals *y*.

27095 An application wishing to check for error situations should set *errno* to 0 before calling  
 27096 *nextafter()*. If *errno* is non-zero on return, or the value NaN is returned, an error has occurred.

27097 **RETURN VALUE**

27098 These functions shall return the next representable double-precision floating-point value  
 27099 following *x* in the direction of *y*. The *nextafter()*, *nextafterf()*, and *nextafterl()* functions shall  
 27100 return *y* if *x* equals *y*.

27101 If *x* or *y* is NaN, then *nextafter()* shall return NaN and may set *errno* to [EDOM].

27102 If *x* is finite and the correct function value would overflow, HUGE\_VAL shall be returned and  
 27103 *errno* set to [ERANGE].

27104 **ERRORS**

27105 These functions shall fail if:

27106 [ERANGE] The correct value would overflow.

27107 These functions may fail if:

27108 [EDOM] The *x* or *y* argument is NaN.

27109 **EXAMPLES**

27110 None.

27111 **APPLICATION USAGE**

27112 None.

27113 **RATIONALE**

27114 None.

27115 **FUTURE DIRECTIONS**

27116 None.

27117 **SEE ALSO**27118 The Base Definitions volume of IEEE Std. 1003.1-200x, `<math.h>`27119 **CHANGE HISTORY**

27120 First released in Issue 4, Version 2.

27121 **Issue 5**

27122 Moved from X/OPEN UNIX extension to BASE.

27123 **Issue 6**27124 The *nextafterf()*, *nextafterl()*, *nexttoward()*, *nexttowardf()*, *nexttowardl()* functions are added for  
27125 alignment with the ISO/IEC 9899:1999 standard.

27126 **NAME**

27127 nftw — walk a file tree

27128 **SYNOPSIS**27129 XSI 

```
#include <ftw.h>
```

```
27130 int nftw(const char *path, int (*fn)(const char *,
27131 const struct stat *, int, struct FTW *), int depth, int flags);
27132
```

27133 **DESCRIPTION**

27134 The *nftw()* function shall recursively descend the directory hierarchy rooted in *path*. The *nftw()*  
 27135 function has a similar effect to *ftw()* except that it takes an additional argument *flags*, which is a  
 27136 bitwise-inclusive OR of zero or more of the following flags:

27137 **FTW\_CHDIR** If set, *nftw()* shall change the current working directory to each directory as it  
 27138 reports files in that directory. If clear, *nftw()* shall not change the current  
 27139 working directory.

27140 **FTW\_DEPTH** If set, *nftw()* shall report all files in a directory before reporting the directory  
 27141 itself. If clear, *nftw()* shall report any directory before reporting the files in that  
 27142 directory.

27143 **FTW\_MOUNT** If set, *nftw()* shall only report files in the same file system as *path*. If clear,  
 27144 *nftw()* shall report all files encountered during the walk.

27145 **FTW\_PHYS** If set, *nftw()* shall perform a physical walk and shall not follow symbolic links.

27146 If **FTW\_PHYS** is clear and **FTW\_DEPTH** is set, *nftw()* shall follow links instead of reporting  
 27147 them, but shall not report any directory that would be a descendant of itself. If **FTW\_PHYS** is  
 27148 clear and **FTW\_DEPTH** is clear, *nftw()* shall follow links instead of reporting them, but shall not  
 27149 report the contents of any directory that would be a descendant of itself.

27150 At each file it encounters, *nftw()* shall call the user-supplied function *fn* with four arguments:

- 27151 • The first argument is the path name of the object.
- 27152 • The second argument is a pointer to the **stat** buffer containing information on the object.
- 27153 • The third argument is an integer giving additional information. Its value is one of the  
 27154 following:

27155 **FTW\_F** The object is a file.

27156 **FTW\_D** The object is a directory.

27157 **FTW\_DP** The object is a directory and subdirectories have been visited. (This condition  
 27158 shall only occur if the **FTW\_DEPTH** flag is included in *flags*.)

27159 **FTW\_SL** The object is a symbolic link. (This condition shall only occur if the **FTW\_PHYS**  
 27160 flag is included in *flags*.)

27161 **FTW\_SLN** The object is a symbolic link that does not name an existing file. (This  
 27162 condition shall only occur if the **FTW\_PHYS** flag is not included in *flags*.)

27163 **FTW\_DNR** The object is a directory that cannot be read. The *fn* function shall not be called  
 27164 for any of its descendants.

27165 **FTW\_NS** The *stat()* function failed on the object because of lack of appropriate  
 27166 permission. The **stat** buffer passed to *fn* is undefined. Failure of *stat()* for any  
 27167 other reason is considered an error and *nftw()* shall return  $-1$ .

27168       • The fourth argument is a pointer to an **FTW** structure. The value of **base** is the offset of the  
 27169       object's file name in the path name passed as the first argument to *fn*. The value of **level**  
 27170       indicates depth relative to the root of the walk, where the root level is 0.

27171       The argument *depth* sets the maximum number of file descriptors that are shall be by *nftw()*  
 27172       while traversing the file tree. At most one file descriptor shall be used for each directory level.

#### 27173 RETURN VALUE

27174       The *nftw()* function shall continue until the first of the following conditions occurs:

- 27175       • An invocation of *fn* shall return a non-zero value, in which case *nftw()* shall return that value.
- 27176       • The *nftw()* function detects an error other than [EACCES] (see FTW\_DNR and FTW\_NS  
 27177       above), in which case *nftw()* shall return  $-1$  and set *errno* to indicate the error.
- 27178       • The tree is exhausted, in which case *nftw()* shall return 0.

#### 27179 ERRORS

27180       The *nftw()* function shall fail if:

27181       [EACCES]       Search permission is denied for any component of *path* or read permission is  
 27182       denied for *path*, or *fn* returns  $-1$  and does not reset *errno*.

27183       [ELOOP]        A loop exists in symbolic links encountered during resolution of the *path*  
 27184       argument.

27185       [ENAMETOOLONG]

27186       The length of the *path* argument exceeds {PATH\_MAX} or a path name  
 27187       component is longer than {NAME\_MAX}.

27188       [ENOENT]        A component of *path* does not name an existing file or *path* is an empty string.

27189       [ENOTDIR]       A component of *path* is not a directory.

27190       The *nftw()* function may fail if:

27191       [ELOOP]        More than {SYMLOOP\_MAX} symbolic links were encountered during  
 27192       resolution of the *path* argument.

27193       [EMFILE]        {OPEN\_MAX} file descriptors are currently open in the calling process.

27194       [ENAMETOOLONG]

27195       Path name resolution of a symbolic link produced an intermediate result  
 27196       whose length exceeds {PATH\_MAX}.

27197       [ENFILE]        Too many files are currently open in the system.

27198       In addition, *errno* may be set if the function pointed by *fn* causes *errno* to be set.

#### 27199 EXAMPLES

27200       The following example walks the **/tmp** directory and its subdirectories, calling the *nftw()*  
 27201       function for every directory entry, to a maximum of 5 levels deep.

```
27202 #include <ftw.h>
```

```
27203 ...
```

```
27204 int nftwfunc(const char *, const struct stat *, int, struct FTW *);
```

```
27205 int nftwfunc(const char *filename, const struct stat *statptr,
27206 int fileflags, struct FTW *pftw)
```

```
27207 {
```

```
27208 return 0;
```

```
27209 }
```

```
27210 ...
27211 char *startpath = "/tmp";
27212 int depth = 5;
27213 int flags = FTW_CHDIR | FTW_DEPTH | FTW_MOUNT;
27214 int ret;

27215 ret = nftw(startpath, nftwfunc, depth, flags);
```

**27216 APPLICATION USAGE**

27217 None.

**27218 RATIONALE**

27219 None.

**27220 FUTURE DIRECTIONS**

27221 None.

**27222 SEE ALSO**

27223 *lstat()*, *opendir()*, *readdir()*, *stat()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ftw.h> |

**27224 CHANGE HISTORY**

27225 First released in Issue 4, Version 2.

**27226 Issue 5**

27227 Moved from X/OPEN UNIX extension to BASE.

27228 In the DESCRIPTION, the definition of the *depth* argument is clarified.

**27229 Issue 6**

27230 The Open Group Base Resolution bwg97-003 is applied. |

27231 The wording of the mandatory [ELOOP] error condition is updated, and a second optional |

27232 [ELOOP] error condition is added. |

27233 **NAME**

27234 nice — change the nice value of a process

27235 **SYNOPSIS**27236 XSI `#include <unistd.h>`27237 `int nice(int incr);`

27238

27239 **DESCRIPTION**

27240 The *nice()* function shall add the value of *incr* to the nice value of the calling process. A process' nice value is a non-negative number for which a more positive value results in less favorable scheduling.

27243 A maximum nice value of  $2*\{NZERO\}-1$  and a minimum nice value of 0 are imposed by the system. Requests for values above or below these limits result in the nice value being set to the corresponding limit. Only a process with appropriate privileges can lower the nice value.

27246 PS|TPS Calling the *nice()* function has no effect on the priority of processes or threads with policy SCHED\_FIFO or SCHED\_RR. The effect on processes or threads with other scheduling policies is implementation-defined.

27249 The nice value set with *nice()* shall be applied to the process. If the process is multi-threaded, the nice value shall affect all system scope threads in the process.

27251 As  $-1$  is a permissible return value in a successful situation, an application wishing to check for error situations should set *errno* to 0, then call *nice()*, and if it returns  $-1$ , check to see if *errno* is non-zero.

27254 **RETURN VALUE**

27255 Upon successful completion, *nice()* shall return the new nice value  $-\{NZERO\}$ . Otherwise,  $-1$  shall be returned, the process' nice value shall not be changed, and *errno* shall be set to indicate the error.

27258 **ERRORS**27259 The *nice()* function shall fail if:

27260 [EPERM] The *incr* argument is negative and the calling process does not have appropriate privileges.

27262 **EXAMPLES**27263 **Changing the Nice Value**

27264 The following example adds the value of the *incr* argument,  $-20$ , to the nice value of the calling process.

27266 `#include <unistd.h>`27267 `...`27268 `int incr = -20;`27269 `int ret;`27270 `ret = nice(incr);`27271 **APPLICATION USAGE**

27272 None.

27273 **RATIONALE**

27274 None.

27275 **FUTURE DIRECTIONS**

27276 None.

27277 **SEE ALSO**

27278 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;limits.h&gt;, &lt;unistd.h&gt;

27279 **CHANGE HISTORY**

27280 First released in Issue 1. Derived from Issue 1 of the SVID.

27281 **Issue 4**

27282 The &lt;unistd.h&gt; header is added to the SYNOPSIS section.

27283 A statement is added to the DESCRIPTION indicating that the nice value can only be lowered by  
27284 a process with appropriate privileges.27285 **Issue 4, Version 2**27286 The RETURN VALUE section is updated for X/OPEN UNIX conformance to define that the  
27287 process' nice value is not changed if an error is detected.27288 **Issue 5**27289 A statement is added to the description indicating the effects of this function on the different  
27290 scheduling policies and multi-threaded processes.

27291 **NAME**

27292 nl\_langinfo — language information

27293 **SYNOPSIS**

27294 XSI #include &lt;langinfo.h&gt;

27295 char \*nl\_langinfo(nl\_item item);

27296

27297 **DESCRIPTION**

27298 The *nl\_langinfo()* function shall return a pointer to a string containing information relevant to  
 27299 the particular language or cultural area defined in the program's locale (see <langinfo.h>). The  
 27300 manifest constant names and values of *item* are defined in <langinfo.h>. For example:

27301 nl\_langinfo(ABDAY\_1)

27302 would return a pointer to the string "Dom" if the identified language was Portuguese, and  
 27303 "Sun" if the identified language was English.

27304 Calls to *setlocale()* with a category corresponding to the category of *item* (see <langinfo.h>), or to  
 27305 the category *LC\_ALL*, may overwrite the array pointed to by the return value.

27306 The *nl\_langinfo()* function need not be reentrant. A function that is not required to be reentrant is  
 27307 not required to be thread-safe.

27308 **RETURN VALUE**

27309 In a locale where *langinfo* data is not defined, *nl\_langinfo()* shall return a pointer to the  
 27310 corresponding string in the POSIX locale. In all locales, *nl\_langinfo()* shall return a pointer to an  
 27311 empty string if *item* contains an invalid setting.

27312 This pointer may point to static data that may be overwritten on the next call.

27313 **ERRORS**

27314 No errors are defined.

27315 **EXAMPLES**27316 **Getting Date and Time Formatting Information**

27317 The following example returns a pointer to a string containing date and time formatting  
 27318 information, as defined in the *LC\_TIME* category of the current locale.

27319 #include &lt;time.h&gt;

27320 #include &lt;langinfo.h&gt;

27321 ...

27322 strftime(datestring, sizeof(datestring), nl\_langinfo(D\_T\_FMT), tm);

27323 ...

27324 **APPLICATION USAGE**

27325 The array pointed to by the return value should not be modified by the program, but may be  
 27326 modified by further calls to *nl\_langinfo()*.

27327 **RATIONALE**

27328 None.

27329 **FUTURE DIRECTIONS**

27330 None.

27331 **SEE ALSO**

27332 *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <langinfo.h>, <nl\_types.h>, the  
27333 Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

27334 **CHANGE HISTORY**

27335 First released in Issue 2.

27336 **Issue 4**

27337 The <nl\_types.h> header is removed from the SYNOPSIS section.

27338 **Issue 5**

27339 The last paragraph of the DESCRIPTION is moved from the APPLICATION USAGE section.

27340 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

27341 **NAME**

27342       nrand48 — generate uniformly distributed pseudo-random non-negative long integers

27343 **SYNOPSIS**

27344 xSI     #include <stdlib.h>

27345       long nrand48(unsigned short xsubi[3]);

27346

27347 **DESCRIPTION**

27348       Refer to *drand48()*.

27349 **NAME**

27350           ntohl — convert values between host and network byte order

27351 **SYNOPSIS**

27352           #include <arpa/inet.h>

27353           uint32\_t ntohl(uint32\_t *netlong*);

27354 **DESCRIPTION**

27355           Refer to *htonl()*.

27356 **NAME**

27357       ntohs — convert values between host and network byte order

27358 **SYNOPSIS**

27359       #include <arpa/inet.h>

27360       uint16\_t ntohs(uint16\_t *netshort*);

27361 **DESCRIPTION**

27362       Refer to *htonl()*.

## 27363 NAME

27364 open — open a file

## 27365 SYNOPSIS

27366 OH #include &lt;sys/stat.h&gt;

27367 #include &lt;fcntl.h&gt;

27368 int open(const char \*path, int oflag, ... );

## 27369 DESCRIPTION

27370 The *open()* function establishes the connection between a file and a file descriptor. It creates an  
 27371 open file description that refers to a file and a file descriptor that refers to that open file  
 27372 description. The file descriptor is used by other I/O functions to refer to that file. The *path*  
 27373 argument points to a path name naming the file.

27374 The *open()* function shall return a file descriptor for the named file that is the lowest file  
 27375 descriptor not currently open for that process. The open file description is new, and therefore the  
 27376 file descriptor does not share it with any other process in the system. The FD\_CLOEXEC file  
 27377 descriptor flag associated with the new file descriptor shall be cleared.

27378 The file offset used to mark the current position within the file shall be set to the beginning of the  
 27379 file.

27380 The file status flags and file access modes of the open file description shall be set according to  
 27381 the value of *oflag*.

27382 Values for *oflag* are constructed by a bitwise-inclusive OR of flags from the following list,  
 27383 defined in <fcntl.h>. Applications shall specify exactly one of the first three values (file access  
 27384 modes) below in the value of *oflag*:

27385 O\_RDONLY Open for reading only.

27386 O\_WRONLY Open for writing only.

27387 O\_RDWR Open for reading and writing. The result is undefined if this flag is applied to  
27388 a FIFO.

27389 Any combination of the following may be used:

27390 O\_APPEND If set, the file offset shall be set to the end of the file prior to each write.

27391 O\_CREAT If the file exists, this flag has no effect except as noted under O\_EXCL below.  
 27392 Otherwise, the file is created; the user ID of the file shall be set to the effective  
 27393 user ID of the process; the group ID of the file shall be set to the group ID of  
 27394 the file's parent directory or to the effective group ID of the process; and the  
 27395 access permission bits (see <sys/stat.h>) of the file mode shall be set to the  
 27396 value of the third argument taken as type **mode\_t** modified as follows: a  
 27397 bitwise AND is performed on the file-mode bits and the corresponding bits in  
 27398 the complement of the process' file mode creation mask. Thus, all bits in the  
 27399 file mode whose corresponding bit in the file mode creation mask is set are  
 27400 cleared. When bits other than the file permission bits are set, the effect is  
 27401 unspecified. The third argument does not affect whether the file is open for  
 27402 reading, writing, or for both.

27403 SIO O\_DSYNC Write I/O operations on the file descriptor complete as defined by  
27404 synchronized I/O data integrity completion

27405 O\_EXCL If O\_CREAT and O\_EXCL are set, *open()* shall fail if the file exists. The check  
 27406 for the existence of the file and the creation of the file if it does not exist shall  
 27407 be atomic with respect to other threads executing *open()* naming the same file

|           |            |                                                                                                                            |
|-----------|------------|----------------------------------------------------------------------------------------------------------------------------|
| 27408     |            | name in the same directory with O_EXCL and O_CREAT set. If O_EXCL and                                                      |
| 27409     |            | O_CREAT are set, and <i>path</i> names a symbolic link, <i>open()</i> shall fail and set                                   |
| 27410     |            | <i>errno</i> to [EEXIST], regardless of the contents of the symbolic link. If O_EXCL is                                    |
| 27411     |            | set and O_CREAT is not set, the result is undefined.                                                                       |
| 27412     | O_NOCTTY   | If set and <i>path</i> identify a terminal device, <i>open()</i> shall not cause the terminal                              |
| 27413     |            | device to become the controlling terminal for the process.                                                                 |
| 27414     | O_NONBLOCK | When opening a FIFO with O_RDONLY or O_WRONLY set:                                                                         |
| 27415     |            | • If O_NONBLOCK is set, an <i>open()</i> for reading-only shall return without                                             |
| 27416     |            | delay. An <i>open()</i> for writing-only shall return an error if no process                                               |
| 27417     |            | currently has the file open for reading.                                                                                   |
| 27418     |            | • If O_NONBLOCK is clear, an <i>open()</i> for reading-only shall block the                                                |
| 27419     |            | calling thread until a thread opens the file for writing. An <i>open()</i> for                                             |
| 27420     |            | writing-only shall block the calling thread until a thread opens the file for                                              |
| 27421     |            | reading.                                                                                                                   |
| 27422     |            | When opening a block special or character special file that supports non-                                                  |
| 27423     |            | blocking opens:                                                                                                            |
| 27424     |            | • If O_NONBLOCK is set, the <i>open()</i> function shall return without blocking                                           |
| 27425     |            | for the device to be ready or available. Subsequent behavior of the device                                                 |
| 27426     |            | is device-specific.                                                                                                        |
| 27427     |            | • If O_NONBLOCK is clear, the <i>open()</i> function shall block the calling                                               |
| 27428     |            | thread until the device is ready or available before returning.                                                            |
| 27429     |            | Otherwise, the behavior of O_NONBLOCK is unspecified.                                                                      |
| 27430 SIO | O_RSYNC    | Read I/O operations on the file descriptor complete at the same level of                                                   |
| 27431     |            | integrity as specified by the O_DSYNC and O_SYNC flags. If both O_DSYNC                                                    |
| 27432     |            | and O_RSYNC are set in <i>oflag</i> , all I/O operations on the file descriptor                                            |
| 27433     |            | complete as defined by synchronized I/O data integrity completion. If both                                                 |
| 27434     |            | O_SYNC and O_RSYNC are set in flags, all I/O operations on the file                                                        |
| 27435     |            | descriptor complete as defined by synchronized I/O file integrity completion.                                              |
| 27436 SIO | O_SYNC     | Write I/O operations on the file descriptor complete as defined by                                                         |
| 27437     |            | synchronized I/O file integrity completion.                                                                                |
| 27438     | O_TRUNC    | If the file exists and is a regular file, and the file is successfully opened                                              |
| 27439     |            | O_RDWR or O_WRONLY, its length is truncated to 0, and the mode and                                                         |
| 27440     |            | owner are unchanged. It shall have no effect on FIFO special files or terminal                                             |
| 27441     |            | device files. Its effect on other file types is implementation-defined. The result                                         |
| 27442     |            | of using O_TRUNC with O_RDONLY is undefined.                                                                               |
| 27443     |            | If O_CREAT is set and the file did not previously exist, upon successful completion, <i>open()</i> shall                   |
| 27444     |            | mark for update the <i>st_atime</i> , <i>st_ctime</i> , and <i>st_mtime</i> fields of the file and the <i>st_ctime</i> and |
| 27445     |            | <i>st_mtime</i> fields of the parent directory.                                                                            |
| 27446     |            | If O_TRUNC is set and the file did previously exist, upon successful completion, <i>open()</i> shall                       |
| 27447     |            | mark for update the <i>st_ctime</i> and <i>st_mtime</i> fields of the file.                                                |
| 27448 SIO |            | If both the O_SYNC and O_DSYNC flags are set, the effect is as if only the O_SYNC flag was set.                            |
| 27449     |            |                                                                                                                            |
| 27450 XSR |            | If <i>path</i> refers to a STREAMS file, <i>oflag</i> may be constructed from O_NONBLOCK OR'ed with                        |
| 27451     |            | either O_RDONLY, O_WRONLY, or O_RDWR. Other flag values are not applicable to                                              |
| 27452     |            | STREAMS devices and have no effect on them. The value O_NONBLOCK affects the operation                                     |

|       |                     |                                                                                                                         |
|-------|---------------------|-------------------------------------------------------------------------------------------------------------------------|
| 27453 |                     | of STREAMS drivers and certain functions applied to file descriptors associated with STREAMS                            |
| 27454 |                     | files. For STREAMS drivers, the implementation of O_NONBLOCK is device-specific.                                        |
| 27455 |                     | If <i>path</i> names the master side of a pseudo-terminal device, then it is unspecified whether <i>open()</i>          |
| 27456 |                     | locks the slave side so that it cannot be opened. Portable applications shall call <i>unlockpt()</i> before             |
| 27457 |                     | opening the slave side.                                                                                                 |
| 27458 |                     | The largest value that can be represented correctly in an object of type <code>off_t</code> shall be established        |
| 27459 |                     | as the offset maximum in the open file description.                                                                     |
| 27460 | <b>RETURN VALUE</b> |                                                                                                                         |
| 27461 |                     | Upon successful completion, the function shall open the file and return a non-negative integer                          |
| 27462 |                     | representing the lowest numbered unused file descriptor. Otherwise, <code>-1</code> shall be returned and               |
| 27463 |                     | <i>errno</i> set to indicate the error. No files shall be created or modified if the function returns <code>-1</code> . |
| 27464 | <b>ERRORS</b>       |                                                                                                                         |
| 27465 |                     | The <i>open()</i> function shall fail if:                                                                               |
| 27466 | [EACCES]            | Search permission is denied on a component of the path prefix, or the file                                              |
| 27467 |                     | exists and the permissions specified by <i>oflag</i> are denied, or the file does not                                   |
| 27468 |                     | exist and write permission is denied for the parent directory of the file to be                                         |
| 27469 |                     | created, or O_TRUNC is specified and write permission is denied.                                                        |
| 27470 | [EEXIST]            | O_CREAT and O_EXCL are set, and the named file exists.                                                                  |
| 27471 | [EINTR]             | A signal was caught during <i>open()</i> .                                                                              |
| 27472 | SIO [EINVAL]        | The implementation does not support synchronized I/O for this file.                                                     |
| 27473 | XSR [EIO]           | The <i>path</i> argument names a STREAMS file and a hangup or error occurred                                            |
| 27474 |                     | during the <i>open()</i> .                                                                                              |
| 27475 | [EISDIR]            | The named file is a directory and <i>oflag</i> includes O_WRONLY or O_RDWR.                                             |
| 27476 | [ELOOP]             | A loop exists in symbolic links encountered during resolution of the <i>path</i>                                        |
| 27477 |                     | argument.                                                                                                               |
| 27478 | [EMFILE]            | {OPEN_MAX} file descriptors are currently open in the calling process.                                                  |
| 27479 | [ENAMETOOLONG]      |                                                                                                                         |
| 27480 |                     | The length of the <i>path</i> argument exceeds {PATH_MAX} or a path name                                                |
| 27481 |                     | component is longer than {NAME_MAX}.                                                                                    |
| 27482 | [ENFILE]            | The maximum allowable number of files is currently open in the system.                                                  |
| 27483 | [ENOENT]            | O_CREAT is not set and the named file does not exist; or O_CREAT is set and                                             |
| 27484 |                     | either the path prefix does not exist or the <i>path</i> argument points to an empty                                    |
| 27485 |                     | string.                                                                                                                 |
| 27486 | XSR [ENOSR]         | The <i>path</i> argument names a STREAMS-based file and the system is unable to                                         |
| 27487 |                     | allocate a STREAM.                                                                                                      |
| 27488 | [ENOSPC]            | The directory or file system that would contain the new file cannot be                                                  |
| 27489 |                     | expanded, the file does not exist, and O_CREAT is specified.                                                            |
| 27490 | [ENOTDIR]           | A component of the path prefix is not a directory.                                                                      |
| 27491 | [ENXIO]             | O_NONBLOCK is set, the named file is a FIFO, O_WRONLY is set, and no                                                    |
| 27492 |                     | process has the file open for reading.                                                                                  |
| 27493 | [ENXIO]             | The named file is a character special or block special file, and the device                                             |
| 27494 |                     | associated with this special file does not exist.                                                                       |

|           |                |                                                                                                                                                                                                                        |
|-----------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 27495     | [EOVERFLOW]    | The named file is a regular file and the size of the file cannot be represented correctly in an object of type <code>off_t</code> .                                                                                    |
| 27496     |                |                                                                                                                                                                                                                        |
| 27497     | [EROFS]        | The named file resides on a read-only file system and either <code>O_WRONLY</code> , <code>O_RDWR</code> , <code>O_CREAT</code> (if file does not exist), or <code>O_TRUNC</code> is set in the <i>oflag</i> argument. |
| 27498     |                |                                                                                                                                                                                                                        |
| 27499     |                |                                                                                                                                                                                                                        |
| 27500     |                | The <code>open()</code> function may fail if:                                                                                                                                                                          |
| 27501 XSR | [EAGAIN]       | The <i>path</i> argument names the slave side of a pseudo-terminal device that is locked.                                                                                                                              |
| 27502     |                |                                                                                                                                                                                                                        |
| 27503     | [EINVAL]       | The value of the <i>oflag</i> argument is not valid.                                                                                                                                                                   |
| 27504     | [ELOOP]        | More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.                                                                                                                 |
| 27505     |                |                                                                                                                                                                                                                        |
| 27506     | [ENAMETOOLONG] |                                                                                                                                                                                                                        |
| 27507     |                | As a result of encountering a symbolic link in resolution of the <i>path</i> argument, the length of the substituted path name string exceeded {PATH_MAX}.                                                             |
| 27508     |                |                                                                                                                                                                                                                        |
| 27509 XSR | [ENOMEM]       | The <i>path</i> argument names a STREAMS file and the system is unable to allocate resources.                                                                                                                          |
| 27510     |                |                                                                                                                                                                                                                        |
| 27511     | [ETXTBSY]      | The file is a pure procedure (shared text) file that is being executed and <i>oflag</i> is <code>O_WRONLY</code> or <code>O_RDWR</code> .                                                                              |
| 27512     |                |                                                                                                                                                                                                                        |

## 27513 EXAMPLES

### 27514 Opening a File for Writing by the Owner

27515 The following example opens the file `/tmp/file`, either by creating it (if it does not already exist),  
 27516 or by truncating its length to 0 (if it does exist). In the former case, if the call creates a new file,  
 27517 the access permission bits in the file mode of the file are set to permit reading and writing by the  
 27518 owner, and to permit reading only by group members and others.

27519 If the call to `open()` is successful, the file is opened for writing.

```
27520 #include <fcntl.h>
27521 ...
27522 int fd;
27523 mode_t mode = S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH;
27524 char *filename = "/tmp/file";
27525 ...
27526 fd = open(filename, O_WRONLY | O_CREAT | O_TRUNC, mode);
27527 ...
```

### 27528 Opening a File Using an Existence Check

27529 The following example uses the `open()` function to try to create the `LOCKFILE` file and open it  
 27530 for writing. Because the `open()` function specifies the `O_EXCL` flag, the call fails if the file already  
 27531 exists. In that case, the program assumes that someone else is updating the password file and  
 27532 exits.

```
27533 #include <fcntl.h>
27534 #include <stdio.h>
27535 #include <stdlib.h>
```

```

27536 #define LOCKFILE "/etc/ptmp"
27537 ...
27538 int pfd; /* Integer for file descriptor returned by open() call. */
27539 ...
27540 if ((pfd = open(LOCKFILE, O_WRONLY | O_CREAT | O_EXCL,
27541 S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH)) == -1)
27542 {
27543 fprintf(stderr, "Cannot open /etc/ptmp. Try again later.\n");
27544 exit(1);
27545 }
27546 ...

```

### 27547 **Opening a File for Writing**

27548 The following example opens a file for writing, creating the file if it does not already exist. If the  
 27549 file does exist, the system truncates the file to zero bytes.

```

27550 #include <fcntl.h>
27551 #include <stdio.h>
27552 #include <stdlib.h>
27553 #define LOCKFILE "/etc/ptmp"
27554 ...
27555 int pfd;
27556 char filename[PATH_MAX+1];
27557 ...
27558 if ((pfd = open(filename, O_WRONLY | O_CREAT | O_TRUNC,
27559 S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH)) == -1)
27560 {
27561 perror("Cannot open output file\n"); exit(1);
27562 }
27563 ...

```

### 27564 **APPLICATION USAGE**

27565 None.

### 27566 **RATIONALE**

27567 Except as specified in this volume of IEEE Std. 1003.1-200x, the flags allowed in *oflag* are not  
 27568 mutually-exclusive and any number of them may be used simultaneously.

27569 Some implementations permit opening FIFOs with O\_RDWR. Since FIFOs could be  
 27570 implemented in other ways, and since two file descriptors can be used to the same effect, this  
 27571 possibility is left as undefined.

27572 See *getgroups()* about the group of a newly created file.

27573 The use of *open()* to create a regular file is preferable to the use of *creat()*, because the latter is  
 27574 redundant and included only for historical reasons.

27575 The use of the O\_TRUNC flag on FIFOs and directories (pipes cannot be *open()*-ed) must be  
 27576 permissible without unexpected side effects (for example, *creat()* on a FIFO must not remove  
 27577 data). Because terminal special files might have type-ahead data stored in the buffer, O\_TRUNC  
 27578 should not affect their content, particularly if a program that normally opens a regular file  
 27579 should open the current controlling terminal instead. Other file types, particularly  
 27580 implementation-defined ones, are left implementation-defined.

- 27581 IEEE Std. 1003.1-200x permits [EACCES] to be returned for conditions other than those explicitly  
27582 listed.
- 27583 The O\_NOCTTY flag was added to allow applications to avoid unintentionally acquiring a  
27584 controlling terminal as a side effect of opening a terminal file. This volume of  
27585 IEEE Std. 1003.1-200x does not specify how a controlling terminal is acquired, but it allows an  
27586 implementation to provide this on *open()* if the O\_NOCTTY flag is not set and other conditions  
27587 specified in the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 11, General Terminal  
27588 Interface are met. The O\_NOCTTY flag is an effective no-op if the file being opened is not a  
27589 terminal device.
- 27590 In historical implementations the value of O\_RDONLY is zero. Because of that, it is not possible  
27591 to detect the presence of O\_RDONLY and another option. Future implementations should  
27592 encode O\_RDONLY and O\_WRONLY as bit flags so that:
- 27593 O\_RDONLY | O\_WRONLY == O\_RDWR
- 27594 In general, the *open()* function follows the symbolic link if *path* names a symbolic link. However,  
27595 the *open()* function, when called with O\_CREAT and O\_EXCL, is required to fail with [EEXIST]  
27596 if *path* names an existing symbolic link, even if the symbolic link refers to a nonexistent file. This  
27597 behavior is required so that privileged applications can create a new file in a known location  
27598 without the possibility that a symbolic link might cause the file to be created in a different  
27599 location.
- 27600 For example, a privileged application that must create a file with a predictable name in a user-  
27601 writable directory, such as the user's home directory, could be compromised if the user creates a  
27602 symbolic link with that name that refers to a nonexistent file in a system directory. If the user can  
27603 influence the contents of a file, the user could compromise the system by creating a new system  
27604 configuration or spool file that would then be interpreted by the system. The test for a symbolic  
27605 link which refers to a nonexistent file must be atomic with the creation of a new file.
- 27606 **FUTURE DIRECTIONS**
- 27607 None.
- 27608 **SEE ALSO**
- 27609 *chmod()*, *close()*, *creat()*, *dup()*, *fcntl()*, *lseek()*, *read()*, *umask()*, *unlockpt()*, *write()*, the Base  
27610 Definitions volume of IEEE Std. 1003.1-200x, <fcntl.h>, <sys/stat.h>, <sys/types.h>
- 27611 **CHANGE HISTORY**
- 27612 First released in Issue 1. Derived from Issue 1 of the SVID.
- 27613 **Issue 4**
- 27614 The <sys/types.h> and <sys/stat.h> headers are now marked as optional (OH); these headers do  
27615 not need to be included on XSI-conformant systems.
- 27616 O\_NDELAY is removed from the list of *oflag* values (this flag was marked WITHDRAWN in  
27617 Issue 3).
- 27618 The [ENXIO] error (for the condition where the file is a character or block special file and the  
27619 associated device does not exist) and the [EINVAL] error are marked as extensions.
- 27620 The following changes are incorporated for alignment with the ISO POSIX-1 standard:
- 27621 • The type of argument *path* is changed from **char\*** to **const char\***.
  - 27622 • Various wording changes are made to the DESCRIPTION to improve clarity and to align the  
27623 text with the ISO POSIX-1 standard.
- 27624 The following changes are incorporated for alignment with the FIPS requirements:

- 27625 • In the DESCRIPTION, the description of O\_CREAT is amended and the relevant part marked  
27626 as an extension.
- 27627 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
27628 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
27629 an extension.
- 27630 **Issue 4, Version 2**
- 27631 The following changes are incorporated for X/OPEN UNIX conformance:
- 27632 • The DESCRIPTION is updated to define the use of open flags with STREAMS files, and to  
27633 identify special considerations when opening the master side of a pseudo-terminal.
- 27634 • In the ERRORS section, the [EIO], [ELOOP], and [ENOSR] mandatory error conditions are  
27635 added, and the [EAGAIN], [ENAMETOOLONG], and [ENOMEM] optional error conditions  
27636 are added.
- 27637 **Issue 5**
- 27638 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
27639 Threads Extension.
- 27640 Large File Summit extensions are added.
- 27641 **Issue 6**
- 27642 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.
- 27643 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:
- 27644 • The [ENAMETOOLONG] error is restored as an error dependent on \_POSIX\_NO\_TRUNC.  
27645 This is since behavior may vary from one file system to another.
- 27646 The following new requirements on POSIX implementations derive from alignment with the  
27647 Single UNIX Specification:
- 27648 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
27649 required for conforming implementations of previous POSIX specifications, it was not  
27650 required for UNIX applications.
- 27651 • In the DESCRIPTION, O\_CREAT is amended to state that the group ID of the file is set to the  
27652 group ID of the file's parent directory or to the effective group ID of the process. This is a  
27653 FIPS requirement.
- 27654 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file  
27655 description. This change is to support large files.
- 27656 • In the ERRORS section, the [E\_OVERFLOW] condition is added. This change is to support  
27657 large files.
- 27658 • The [ENXIO] mandatory error condition is added.
- 27659 • The [EINVAL], [ENAMETOOLONG], and [ETXTBSY] optional error conditions are added.
- 27660 The DESCRIPTION and ERRORS sections are updated so that items related to the optional XSI  
27661 STREAMS Option Group are marked.
- 27662 The following changes were made to align with the IEEE P1003.1a draft standard:
- 27663 • An explanation is added of the effect of the O\_CREAT and O\_EXCL flags when the path  
27664 refers to a symbolic link.
- 27665 • The [ELOOP] optional error condition is added.

27666 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

27667 The DESCRIPTION of O\_EXCL is updated in response to IEEE PASC Interpretation 1003.1c #48. |

27668 **NAME**

27669            opendir — open a directory

27670 **SYNOPSIS**

27671            #include &lt;dirent.h&gt;

27672            DIR \*opendir(const char \*dirname);

27673 **DESCRIPTION**

27674            The *opendir()* function shall open a directory stream corresponding to the directory named by  
 27675            the *dirname* argument. The directory stream is positioned at the first entry. If the type **DIR** is  
 27676            implemented using a file descriptor, applications shall only be able to open up to a total of  
 27677            {OPEN\_MAX} files and directories.

27678 **RETURN VALUE**

27679            Upon successful completion, *opendir()* shall return a pointer to an object of type **DIR**.  
 27680            Otherwise, a null pointer shall be returned and *errno* set to indicate the error.

27681 **ERRORS**27682            The *opendir()* function shall fail if:

27683            [EACCES]            Search permission is denied for the component of the path prefix of *dirname* or  
 27684            read permission is denied for *dirname*.

27685            [ELOOP]            A loop exists in symbolic links encountered during resolution of the *dirname*  
 27686            argument.

27687            [ENAMETOOLONG]       The length of the *dirname* argument exceeds {PATH\_MAX} or a path name  
 27688            component is longer than {NAME\_MAX}.

27690            [ENOENT]           A component of *dirname* does not name an existing directory or *dirname* is an  
 27691            empty string.

27692            [ENOTDIR]          A component of *dirname* is not a directory.

27693            The *opendir()* function may fail if:

27694            [ELOOP]            More than {SYMLOOP\_MAX} symbolic links were encountered during  
 27695            resolution of the *dirname* argument.

27696            [EMFILE]           {OPEN\_MAX} file descriptors are currently open in the calling process.

27697            [ENAMETOOLONG]       As a result of encountering a symbolic link in resolution of the *dirname*  
 27698            argument, the length of the substituted path name string exceeded  
 27699            {PATH\_MAX}.

27701            [ENFILE]          Too many files are currently open in the system.

27702 **EXAMPLES**

27703 **Open a Directory Stream**

27704 The following program fragment demonstrates how the *opendir()* function is used.

```

27705 #include <sys/types.h>
27706 #include <dirent.h>
27707 #include <libgen.h>
27708 ...
27709 DIR *dir;
27710 struct dirent *dp;
27711 ...
27712 if ((dir = opendir(".")) == NULL) {
27713 perror("Cannot open .");
27714 exit(1);
27715 }
27716 while ((dp = readdir(dir)) != NULL) {
27717 ...

```

27718 **APPLICATION USAGE**

27719 The *opendir()* function should be used in conjunction with *readdir()*, *closedir()*, and *rewinddir()* to  
 27720 examine the contents of the directory (see the EXAMPLES section in *readdir()*). This method is  
 27721 recommended for portability.

27722 **RATIONALE**

27723 Based on historical implementations, the rules about file descriptors apply to directory streams  
 27724 as well. However, this volume of IEEE Std. 1003.1-200x does not mandate that the directory  
 27725 stream be implemented using file descriptors. The description of *closedir()* clarifies that if a file  
 27726 descriptor is used for the directory stream, it is mandatory that *closedir()* deallocate the file  
 27727 descriptor. When a file descriptor is used to implement the directory stream, it behaves as if the  
 27728 FD\_CLOEXEC had been set for the file descriptor.

27729 The directory entries for dot and dot-dot are optional. This volume of IEEE Std. 1003.1-200x does  
 27730 not provide a way to test *a priori* for their existence because an application that is portable must  
 27731 be written to look for (and usually ignore) those entries. Writing code that presumes that they  
 27732 are the first two entries does not always work, as many implementations permit them to be  
 27733 other than the first two entries, with a “normal” entry preceding them. There is negligible value  
 27734 in providing a way to determine what the implementation does because the code to deal with  
 27735 dot and dot-dot must be written in any case and because such a flag would add to the list of  
 27736 those flags (which has proven in itself to be objectionable) and might be abused.

27737 Since the structure and buffer allocation, if any, for directory operations are defined by the  
 27738 implementation, this volume of IEEE Std. 1003.1-200x imposes no portability requirements for  
 27739 erroneous program constructs, erroneous data, or the use of indeterminate values such as the  
 27740 use or referencing of a *dirp* value or a **dirent** structure value after a directory stream has been  
 27741 closed or after a *fork()* or one of the *exec* function calls.

27742 **FUTURE DIRECTIONS**

27743 None.

27744 **SEE ALSO**

27745 *closedir()*, *lstat()*, *readdir()*, *rewinddir()*, *symlink()*, the Base Definitions volume of  
 27746 IEEE Std. 1003.1-200x, **<dirent.h>**, **<limits.h>**, **<sys/types.h>**

27747 **CHANGE HISTORY**

27748 First released in Issue 2.

27749 **Issue 4**27750 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on  
27751 XSI-conformant systems.27752 In the DESCRIPTION, the following sentence is moved to the Base Definitions volume of  
27753 IEEE Std. 1003.1-200x: “The type **DIR**, which is defined in `<dirent.h>`, represents a *directory*  
27754 *stream*, which is an ordered sequence of all directory entries in a particular directory.”

27755 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 27756
- The type of argument *dirname* is changed from `char*` to `const char*`.
  - The generation of an [ENOENT] error when *dirname* points to an empty string is made  
27757 mandatory.  
27758

27759 The following change is incorporated for alignment with the FIPS requirements:

- 27760
- In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
27761 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
27762 an extension.

27763 **Issue 4, Version 2**

27764 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 27765
- It states that [ELOOP] is returned if too many symbolic links are encountered during path  
27766 name resolution.
  - A second [ENAMETOOLONG] condition is defined that may report excessive length of an  
27767 intermediate result of path name resolution of a symbolic link.  
27768

27769 **Issue 6**27770 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

27771 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 27772
- The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.  
27773 This is since behavior may vary from one file system to another.

27774 The following new requirements on POSIX implementations derive from alignment with the  
27775 Single UNIX Specification:

- 27776
- The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
27777 required for conforming implementations of previous POSIX specifications, it was not  
27778 required for UNIX applications.
  - The [ELOOP] mandatory error condition is added.  
27779
  - A second [ENAMETOOLONG] is added as an optional error condition.  
27780

27781 The following changes were made to align with the IEEE P1003.1a draft standard:

- 27782
- The [ELOOP] optional error condition is added.

27783 **NAME**

27784            **openlog** — open a connection to the logging facility

27785 **SYNOPSIS**

27786 XSI        #include <syslog.h>

27787            void **openlog**(const char \**ident*, int *logopt*, int *facility*);

27788

27789 **DESCRIPTION**

27790            Refer to *closelog*().

27791 **NAME**

27792           optarg, opterr, optind, optopt — options parsing variables

27793 **SYNOPSIS**

27794           #include <unistd.h>

27795           extern char \*optarg;

27796           extern int opterr, optind, optopt;

27797 **DESCRIPTION**

27798           Refer to *getopt()*.

27799 **NAME**

27800 pathconf — get configurable path name variables

27801 **SYNOPSIS**

27802 #include <unistd.h>

27803 long pathconf(const char \*path, int name);

27804 **DESCRIPTION**

27805 Refer to *fpathconf()*.

27806 **NAME**

27807 `pause` — suspend the thread until a signal is received

27808 **SYNOPSIS**

27809 `#include <unistd.h>`

27810 `int pause(void);`

27811 **DESCRIPTION**

27812 The `pause()` function shall suspend the calling thread until delivery of a signal whose action is  
27813 either to execute a signal-catching function or to terminate the process.

27814 If the action is to terminate the process, `pause()` shall not return.

27815 If the action is to execute a signal-catching function, `pause()` shall return after the signal-catching  
27816 function returns.

27817 **RETURN VALUE**

27818 Since `pause()` suspends thread execution indefinitely unless interrupted by a signal, there is no  
27819 successful completion return value. A value of `-1` shall be returned and `errno` set to indicate the  
27820 error.

27821 **ERRORS**

27822 The `pause()` function shall fail if:

27823 [EINTR] A signal is caught by the calling process and control is returned from the  
27824 signal-catching function.

27825 **EXAMPLES**

27826 None.

27827 **APPLICATION USAGE**

27828 Many common uses of `pause()` have timing windows. The scenario involves checking a  
27829 condition related to a signal and, if the signal has not occurred, calling `pause()`. When the signal  
27830 occurs between the check and the call to `pause()`, the process often blocks indefinitely. The  
27831 `sigprocmask()` and `sigsuspend()` functions can be used to avoid this type of problem.

27832 **RATIONALE**

27833 None.

27834 **FUTURE DIRECTIONS**

27835 None.

27836 **SEE ALSO**

27837 `sigsuspend()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<unistd.h>`

27838 **CHANGE HISTORY**

27839 First released in Issue 1. Derived from Issue 1 of the SVID.

27840 **Issue 4**

27841 The `<unistd.h>` header is added to the SYNOPSIS section.

27842 In the RETURN VALUE section, the text is expanded to indicate that process execution is  
27843 suspended indefinitely “unless interrupted by a signal”.

27844 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 27845 • The argument list is explicitly defined as **void**.

27846 **Issue 5**

27847 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

27848 **Issue 6**

27849 The APPLICATION USAGE section is added.

27850 **NAME**

27851 pclose — close a pipe stream to or from a process

27852 **SYNOPSIS**

27853 #include &lt;stdio.h&gt;

27854 int pclose(FILE \*stream);

27855 **DESCRIPTION**

27856 The *pclose()* function shall close a stream that was opened by *popen()*, wait for the command to  
 27857 terminate, and return the termination status of the process that was running the command  
 27858 language interpreter. However, if a call caused the termination status to be unavailable to  
 27859 *pclose()*, then *pclose()* shall return  $-1$  with *errno* set to [ECHILD] to report this situation. This can  
 27860 happen if the application calls one of the following functions:

- 27861 • *wait()*
- 27862 • *waitpid()* with a *pid* argument less than or equal to 0 or equal to the process ID of the  
 27863 command line interpreter
- 27864 • Any other function not defined in this volume of IEEE Std. 1003.1-200x that could do one of  
 27865 the above

27866 In any case, *pclose()* shall not return before the child process created by *popen()* has terminated.

27867 If the command language interpreter cannot be executed, the child termination status returned  
 27868 by *pclose()* is as if the command language interpreter terminated using *exit(127)* or *\_exit(127)*.

27869 The *pclose()* function shall not affect the termination status of any child of the calling process  
 27870 other than the one created by *popen()* for the associated stream.

27871 If the argument *stream* to *pclose()* is not a pointer to a stream created by *popen()*, the result of  
 27872 *pclose()* is undefined.

27873 **RETURN VALUE**

27874 Upon successful return, *pclose()* shall return the termination status of the command language  
 27875 interpreter. Otherwise, *pclose()* shall return  $-1$  and set *errno* to indicate the error.

27876 **ERRORS**

27877 The *pclose()* function shall fail if:

27878 [ECHILD] The status of the child process could not be obtained, as described above.

27879 **EXAMPLES**

27880 None.

27881 **APPLICATION USAGE**

27882 None.

27883 **RATIONALE**

27884 There is a requirement that *pclose()* not return before the child process terminates. This is  
 27885 intended to disallow implementations that return [EINTR] if a signal is received while waiting.  
 27886 If *pclose()* returned before the child terminated, there would be no way for the application to  
 27887 discover which child used to be associated with the stream, and it could not do the cleanup  
 27888 itself.

27889 If the stream pointed to by *stream* was not created by *popen()*, historical implementations of  
 27890 *pclose()* return  $-1$  without setting *errno*. To avoid requiring *pclose()* to set *errno* in this case,  
 27891 IEEE Std. 1003.1-200x makes the behavior unspecified. An application should not use *pclose()* to  
 27892 close any stream that was not created by *popen()*.

27893 Some historical implementations of *pclose()* either block or ignore the signals SIGINT, SIGQUIT,  
 27894 and SIGHUP while waiting for the child process to terminate. Since this behavior is not  
 27895 described for the *pclose()* function in IEEE Std. 1003.1-200x, such implementations are not  
 27896 conforming. Also, some historical implementations return [EINTR] if a signal is received, even  
 27897 though the child process has not terminated. Such implementations are also considered non-  
 27898 conforming.

27899 Consider, for example, an application that uses:

```
27900 popen("command", "r")
```

27901 to start *command*, which is part of the same application. The parent writes a prompt to its  
 27902 standard output (presumably the terminal) and then reads from the stream. The child reads the  
 27903 response from the user, does some transformation on the response (path name expansion,  
 27904 perhaps) and writes the result to its standard output. The parent process reads the result from  
 27905 the pipe, does something with it, and prints another prompt. The cycle repeats. Assuming that  
 27906 both processes do appropriate buffer flushing, this would be expected to work.

27907 To conform to IEEE Std. 1003.1-200x, *pclose()* must use *waitpid()*, or some similar function,  
 27908 instead of *wait()*.

27909 The code sample below illustrates how the *pclose()* function might be implemented on a system  
 27910 conforming to IEEE Std. 1003.1-200x.

```
27911 int pclose(FILE *stream)
27912 {
27913 int stat;
27914 pid_t pid;
27915
27916 pid = <pid for process created for stream by popen(>
27917 (void) fclose(stream);
27918 while (waitpid(pid, &stat, 0) == -1) {
27919 if (errno != EINTR){
27920 stat = -1;
27921 break;
27922 }
27923 }
27924 return(stat);
27925 }
```

#### 27925 FUTURE DIRECTIONS

27926 None.

#### 27927 SEE ALSO

27928 *fork()*, *popen()*, *waitpid()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdio.h>`

#### 27929 CHANGE HISTORY

27930 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 27931 Issue 4

27932 The following changes are incorporated for alignment with the ISO POSIX-2 standard:

- 27933 • The function is no longer marked as an extension.
- 27934 • The simple DESCRIPTION given in Issue 3 is replaced with a more complete description in  
 27935 this issue. In particular, interactions between this function and *wait()* and *waitpid()* are  
 27936 defined.

27937 **NAME**

27938           perror — write error messages to standard error

27939 **SYNOPSIS**

27940           #include &lt;stdio.h&gt;

27941           void perror(const char \*s);

27942 **DESCRIPTION**

27943 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
 27944       conflict between the requirements described here and the ISO C standard is unintentional. This  
 27945       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

27946       The *perror()* function maps the error number accessed through the symbol *errno* to a language-  
 27947       dependent error message, which shall be written to the standard error stream as follows:

- 27948           • First (if *s* is not a null pointer and the character pointed to by *s* is not the null byte), the string  
 27949           pointed to by *s* followed by a colon and a <space> character.
- 27950           • Then an error message string followed by a <newline> character.

27951       The contents of the error message strings are the same as those returned by *strerror()* with  
 27952       argument *errno*.

27953 cx       The *perror()* function shall mark the file associated with the standard error stream as having  
 27954       been written (*st\_ctime*, *st\_mtime* marked for update) at some time between its successful  
 27955       completion and *exit()*, *abort()*, or the completion of *fflush()* or *fclose()* on *stderr*.

27956       The *perror()* function shall not change the orientation of the standard error stream.

27957 **RETURN VALUE**27958           The *perror()* function shall return no value.27959 **ERRORS**

27960           No errors are defined.

27961 **EXAMPLES**27962           **Printing an Error Message for a Function**

27963       The following example replaces *bufptr* with a buffer that is the necessary size. If an error occurs,  
 27964       the *perror()* function prints a message and the program exits.

```

27965 #include <stdio.h>
27966 #include <stdlib.h>
27967 ...
27968 char *bufptr;
27969 size_t szbuf;
27970 ...
27971 if ((bufptr = malloc(szbuf)) == NULL) {
27972 perror("malloc"); exit(2);
27973 }
27974 ...
```

27975 **APPLICATION USAGE**

27976           None.

27977 **RATIONALE**

27978 None.

27979 **FUTURE DIRECTIONS**

27980 None.

27981 **SEE ALSO**27982 *strerror()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdio.h**>27983 **CHANGE HISTORY**

27984 First released in Issue 1. Derived from Issue 1 of the SVID.

27985 **Issue 4**27986 The language for error message strings was given as implementation-defined in Issue 3. In this  
27987 issue, they are defined as language-dependent.

27988 The following change is incorporated for alignment with the ISO POSIX-1 standard:

27989 • A paragraph is added to the DESCRIPTION defining the effects of this function on the  
27990 *st\_ctime* and *st\_mtime* fields of the standard error stream.

27991 The following change is incorporated for alignment with the ISO C standard:

27992 • The type of argument *s* is changed from **char\*** to **const char\***.27993 **Issue 5**27994 A paragraph is added to the DESCRIPTION indicating that *perror()* does not change the  
27995 orientation of the standard error stream.27996 **Issue 6**

27997 Extensions beyond the ISO C standard are now marked.

27998 **NAME**

27999 pipe — create an interprocess channel

28000 **SYNOPSIS**

28001 #include &lt;unistd.h&gt;

28002 int pipe(int *filides*[2]);28003 **DESCRIPTION**

28004 The *pipe()* function shall create a pipe and place two file descriptors, one each into the  
 28005 arguments *filides*[0] and *filides*[1], that refer to the open file descriptions for the read and write  
 28006 ends of the pipe. Their integer values shall be the two lowest available at the time of the *pipe()*  
 28007 call. The O\_NONBLOCK and FD\_CLOEXEC flags shall be clear on both file descriptors. (The  
 28008 *fcntl()* function can be used to set both these flags.)

28009 Data can be written to the file descriptor *filides*[1] and read from the file descriptor *filides*[0]. A  
 28010 read on the file descriptor *filides*[0] shall access data written to the file descriptor *filides*[1] on a  
 28011 first-in-first-out basis. It is unspecified whether *filides*[0] is also open for writing and whether  
 28012 *filides*[1] is also open for reading.

28013 A process has the pipe open for reading (correspondingly writing) if it has a file descriptor open  
 28014 that refers to the read end, *filides*[0] (write end, *filides*[1]).

28015 Upon successful completion, *pipe()* shall mark for update the *st\_atime*, *st\_ctime*, and *st\_mtime*  
 28016 fields of the pipe.

28017 **RETURN VALUE**

28018 Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to  
 28019 indicate the error.

28020 **ERRORS**28021 The *pipe()* function shall fail if:

28022 [EMFILE] More than {OPEN\_MAX} minus two file descriptors are already in use by this  
 28023 process.

28024 [ENFILE] The number of simultaneously open files in the system would exceed a  
 28025 system-imposed limit.

28026 **EXAMPLES**

28027 None.

28028 **APPLICATION USAGE**

28029 None.

28030 **RATIONALE**

28031 The wording carefully avoids using the verb “to open” in order to avoid any implication of use  
 28032 of *open()*; see also *write()*.

28033 **FUTURE DIRECTIONS**

28034 None.

28035 **SEE ALSO**

28036 *fcntl()*, *read()*, *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <*fcntl.h*>,  
 28037 <*unistd.h*>

28038 **CHANGE HISTORY**

28039 First released in Issue 1. Derived from Issue 1 of the SVID.

28040 **Issue 4**

28041       The <**unistd.h**> header is added to the SYNOPSIS section.

28042 **Issue 4, Version 2**

28043       The DESCRIPTION is updated for X/OPEN UNIX conformance to indicate that certain  
28044       dispositions of *fildev[0]* and *fildev[1]* are unspecified.

28045 **Issue 6**

28046       The following new requirements on POSIX implementations derive from alignment with the  
28047       Single UNIX Specification:

- 28048       • The DESCRIPTION is updated to indicate that certain dispositions of *fildev[0]* and *fildev[1]*  
28049       are unspecified.

## 28050 NAME

28051 poll — input/output multiplexing

## 28052 SYNOPSIS

28053 XSI #include &lt;poll.h&gt;

28054 int poll(struct pollfd *fds*[], nfd\_t *nfd*s, int *timeout*);

28055

## 28056 DESCRIPTION

28057 The *poll()* function provides applications with a mechanism for multiplexing input/output over  
 28058 a set of file descriptors. For each member of the array pointed to by *fds*, *poll()* examines the given  
 28059 file descriptor for the event(s) specified in *events*. The number of **pollfd** structures in the *fds*  
 28060 array is specified by *nfd*s. The *poll()* function identifies those file descriptors on which an  
 28061 application can read or write data, or on which certain events have occurred.

28062 The *fds* argument specifies the file descriptors to be examined and the events of interest for each  
 28063 file descriptor. It is a pointer to an array with one member for each open file descriptor of  
 28064 interest. The array's members are **pollfd** structures within which *fd* specifies an open file  
 28065 descriptor and *events* and *revents* are bitmasks constructed by OR'ing a combination of the  
 28066 following event flags:

28067 XSR POLLIN Data other than high-priority data may be read without blocking. For  
 28068 STREAMS, this flag is set in *revents* even if the message is of zero length.

28069 XSR POLLRDNORM Normal data (priority band equals 0) may be read without blocking. For  
 28070 STREAMS, this flag is set in *revents* even if the message is of zero length.

28071 XSR POLLRDBAND Data from a non-zero priority band may be read without blocking. For  
 28072 STREAMS, this flag is set in *revents* even if the message is of zero length.

28073 XSR POLLPRI High-priority data may be received without blocking. For STREAMS, this flag  
 28074 is set in *revents* even if the message is of zero length.

28075 POLLOUT Normal data (priority band equals 0) may be written without blocking.

28076 POLLWRNORM Same as POLLOUT.

28077 POLLWRBAND Priority data (priority band >0) may be written. This event only examines  
 28078 bands that have been written to at least once.

28079 POLLERR An error has occurred on the device or stream. This flag is only valid in the  
 28080 *revents* bitmask; it is ignored in the *events* member.

28081 POLLHUP The device has been disconnected. This event and POLLOUT are mutually-  
 28082 exclusive; a stream can never be writable if a hangup has occurred. However,  
 28083 this event and POLLIN, POLLRDNORM, POLLRDBAND, or POLLPRI are not  
 28084 mutually-exclusive. This flag is only valid in the *revents* bitmask; it is ignored  
 28085 in the *events* member.

28086 POLLNVAL The specified *fd* value is invalid. This flag is only valid in the *revents* member;  
 28087 it is ignored in the *events* member.

28088 If the value of *fd* is less than 0, *events* is ignored, and *revents* is set to 0 in that entry on return from  
 28089 *poll()*.

28090 In each **pollfd** structure, *poll()* clears the *revents* member, except that where the application  
 28091 requested a report on a condition by setting one of the bits of *events* listed above, *poll()* sets the  
 28092 corresponding bit in *revents* if the requested condition is true. In addition, *poll()* sets the  
 28093 POLLHUP, POLLERR, and POLLNVAL flag in *revents* if the condition is true, even if the

- 28094 application did not set the corresponding bit in *events*.
- 28095 If none of the defined events have occurred on any selected file descriptor, *poll()* waits at least  
28096 *timeout* milliseconds for an event to occur on any of the selected file descriptors. If the value of  
28097 *timeout* is 0, *poll()* shall return immediately. If the value of *timeout* is -1, *poll()* shall block until a  
28098 requested event occurs or until the call is interrupted.
- 28099 Implementations may place limitations on the granularity of timeout intervals. If the requested  
28100 timeout interval requires a finer granularity than the implementation supports, the actual  
28101 timeout interval shall be rounded up to the next supported value.
- 28102 The *poll()* function is not affected by the O\_NONBLOCK flag.
- 28103 XSR The *poll()* function supports regular files, terminal and pseudo-terminal devices, STREAMS-  
28104 based files, FIFOs, pipes, and sockets. The behavior of *poll()* on elements of *fds* that refer to other  
28105 types of file is unspecified.
- 28106 Regular files always poll TRUE for reading and writing.
- 28107 The *poll()* function supports sockets.
- 28108 A file descriptor for a socket that is listening for connections shall indicate that it is ready for  
28109 reading, once connections are available. A file descriptor for a socket that is connecting  
28110 asynchronously shall indicate that it is ready for writing, once a connection has been established.
- 28111 **RETURN VALUE**
- 28112 Upon successful completion, *poll()* shall return a non-negative value. A positive value indicates  
28113 the total number of file descriptors that have been selected (that is, file descriptors for which the  
28114 *revents* member is non-zero). A value of 0 indicates that the call timed out and no file descriptors  
28115 have been selected. Upon failure, *poll()* shall return -1 and set *errno* to indicate the error.
- 28116 **ERRORS**
- 28117 The *poll()* function shall fail if:
- 28118 [EAGAIN] The allocation of internal data structures failed but a subsequent request may  
28119 succeed.
- 28120 [EINTR] A signal was caught during *poll()*.
- 28121 XSR [EINVAL] The *nfds* argument is greater than {OPEN\_MAX}, or one of the *fd* members  
28122 refers to a STREAM or multiplexer that is linked (directly or indirectly)  
28123 downstream from a multiplexer.
- 28124 **EXAMPLES**
- 28125 **Checking for Events on a Stream**
- 28126 The following example opens a pair of STREAMS devices and then waits for either one to  
28127 become writable. This example proceeds as follows:
- 28128 1. Sets the *timeout* parameter to 500 milliseconds.
  - 28129 2. Opens the STREAMS devices */dev/dev0* and */dev/dev1*, and then polls them, specifying  
28130 POLLOUT and POLLWRBAND as the events of interest.
- 28131 The STREAMS device names */dev/dev0* and */dev/dev1* are only examples of how  
28132 STREAMS devices can be named; STREAMS naming conventions may vary among  
28133 systems conforming to the IEEE Std. 1003.1-200x.
- 28134 3. Uses the *ret* variable to determine whether an event has occurred on either of the two  
28135 STREAMS. The *poll()* function is given 500 milliseconds to wait for an event to occur (if it

- 28136 has not occurred prior to the *poll()* call).
- 28137 4. Checks the returned value of *ret*. If a positive value is returned, one of the following can  
28138 be done:
- 28139 a. Priority data can be written to the open STREAM on priority bands greater than 0,  
28140 because the POLLWRBAND event occurred on the open STREAM (*fds*[0] or *fds*[1]).
  - 28141 b. Data can be written to the open STREAM on priority-band 0, because the POLLOUT  
28142 event occurred on the open STREAM (*fds*[0] or *fds*[1]).
- 28143 5. If the returned value is not a positive value, permission to write data to the open STREAM  
28144 (on any priority band) is denied.
- 28145 6. If the POLLHUP event occurs on the open STREAM (*fds*[0] or *fds*[1]), the device on the  
28146 open STREAM has disconnected.

```

28147 #include <stropts.h>
28148 #include <poll.h>
28149 ...
28150 struct pollfd fds[2];
28151 int timeout_msecs = 500;
28152 int ret;
28153 int i;

28154 /* Open STREAMS device. */
28155 fds[0].fd = open("/dev/dev0", ...);
28156 fds[1].fd = open("/dev/dev1", ...);
28157 fds[0].events = POLLOUT | POLLWRBAND;
28158 fds[1].events = POLLOUT | POLLWRBAND;

28159 ret = poll(fds, 2, timeout_msecs);

28160 if (ret > 0) {
28161 /* An event on one of the fds has occurred. */
28162 for (i=0; i<2; i++) {
28163 if (fds[i].revents & POLLWRBAND) {
28164 /* Priority data may be written on device number i. */
28165 ...
28166 }
28167 if (fds[i].revents & POLLOUT) {
28168 /* Data may be written on device number i. */
28169 ...
28170 }
28171 if (fds[i].revents & POLLHUP) {
28172 /* A hangup has occurred on device number i. */
28173 ...
28174 }
28175 }
28176 }

```

#### 28177 APPLICATION USAGE

28178 None.

28179 **RATIONALE**

28180 None.

28181 **FUTURE DIRECTIONS**

28182 None.

28183 **SEE ALSO**28184 *getmsg()*, *putmsg()*, *read()*, *select()*, *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |28185 `<poll.h>`, `<stropts.h>`, Section 2.6 (on page 539)28186 **CHANGE HISTORY**

28187 First released in Issue 4, Version 2.

28188 **Issue 5**

28189 Moved from X/OPEN UNIX extension to BASE.

28190 The description of POLLWRBAND is updated.

28191 **Issue 6**

28192 Text referring to sockets is added to the DESCRIPTION.

28193 Text relating to the XSI STREAMS Option Group is marked.

28194 **NAME**

28195 popen — initiate pipe streams to or from a process

28196 **SYNOPSIS**

28197 #include &lt;stdio.h&gt;

28198 FILE \*popen(const char \*command, const char \*mode);

28199 **DESCRIPTION**

28200 The *popen()* function shall execute the command specified by the string *command*. It creates a  
 28201 pipe between the calling program and the executed command, and returns a pointer to a stream  
 28202 that can be used to either read from or write to the pipe.

28203 The environment of the executed command shall be as if a child process were created within the  
 28204 *popen()* call using the *fork()* function, and the child used *execl()* to invoke a command line  
 28205 interpreter.

28206 If the implementation supports the Shell and Utilities volume of IEEE Std. 1003.1-200x, the  
 28207 environment of the executed command is as if a child process were created within the *popen()*  
 28208 call using *fork()*, and the child invoked the *sh* utility using the call:

28209 `execl(shell_path, "sh", "-c", command, (char *)0);`28210 where *shell\_path* is an unspecified path name for the *sh* utility.

28211 The *popen()* function ensures that any streams from previous *popen()* calls that remain open in  
 28212 the parent process are closed in the new child process.

28213 The *mode* argument to *popen()* is a string that specifies I/O mode:

- 28214 1. If *mode* is *r*, when the child process is started, its file descriptor `STDOUT_FILENO` shall be  
 28215 the writable end of the pipe, and the file descriptor *fileno(stream)* in the calling process,  
 28216 where *stream* is the stream pointer returned by *popen()*, shall be the readable end of the  
 28217 pipe.
- 28218 2. If *mode* is *w*, when the child process is started its file descriptor `STDIN_FILENO` shall be  
 28219 the readable end of the pipe, and the file descriptor *fileno(stream)* in the calling process,  
 28220 where *stream* is the stream pointer returned by *popen()*, shall be the writable end of the  
 28221 pipe.
- 28222 3. If *mode* is any other value, the result is undefined.

28223 After *popen()*, both the parent and the child process shall be capable of executing independently  
 28224 before either terminates.

28225 Pipe streams are byte-oriented.

28226 **RETURN VALUE**

28227 Upon successful completion, *popen()* shall return a pointer to an open stream that can be used to  
 28228 read or write to the pipe. Otherwise, it shall return a null pointer and may set *errno* to indicate  
 28229 the error.

28230 **ERRORS**28231 The *popen()* function may fail if:

28232 [EMFILE] {FOPEN\_MAX} or {STREAM\_MAX} streams are currently open in the calling  
 28233 process.

28234 [EINVAL] The *mode* argument is invalid.28235 The *popen()* function may also set *errno* values as described by *fork()* or *pipe()*.

28236 **EXAMPLES**

28237 None.

28238 **APPLICATION USAGE**

28239 Because open files are shared, a mode *r* command can be used as an input filter and a mode *w*  
28240 command as an output filter.

28241 Buffered reading before opening an input filter may leave the standard input of that filter  
28242 mispositioned. Similar problems with an output filter may be prevented by careful buffer  
28243 flushing; for example, with *fflush()*.

28244 A stream opened by *popen()* should be closed by *pclose()*.

28245 The behavior of *popen()* is specified for values of *mode* of *r* and *w*. Other modes such as *rb* and  
28246 *wb* might be supported by specific implementations, but these would not be portable features.  
28247 Note that historical implementations of *popen()* only check to see if the first character of *mode* is  
28248 *r*. Thus, a *mode* of *robert the robot* would be treated as *mode r*, and a *mode* of *anything else* would be  
28249 treated as *mode w*.

28250 If the application calls *waitpid()* or *waitid()* with a *pid* argument greater than 0, and it still has a  
28251 stream that was called with *popen()* open, it must ensure that *pid* does not refer to the process  
28252 started by *popen()*.

28253 To determine whether or not the environment specified in the Shell and Utilities volume of  
28254 IEEE Std. 1003.1-200x is present, use the function call:

28255 `sysconf(_SC_2_VERSION)`

28256 (See *sysconf()*).

28257 **RATIONALE**

28258 The *popen()* function should not be used by programs that have set user (or group) ID privileges.  
28259 The *fork()* and *exec* family of functions (except *execlp()* and *execvp()*), should be used instead.  
28260 This prevents any unforeseen manipulation of the environment of the user that could cause  
28261 execution of commands not anticipated by the calling program.

28262 If the original and *popen()*ed processes both intend to read or write or read and write a common  
28263 file, and either will be using FILE-type C functions (*fread()*, *fwrite()*, and so on), the rules for  
28264 sharing file handles must be observed (see Section 2.5.1 (on page 535)).

28265 Because open files are shared, a mode *r* argument can be used as an input filter and a mode *w*  
28266 argument as an output filter.

28267 The behavior of *popen()* is specified for modes of *r* and *w*. Other modes such as *rb* and *wb* might  
28268 be supported by specific implementations, but these would not be portable features. Note that  
28269 historical implementations of *popen()* only check to see if the first character of *mode* is '*r*'.  
28270 Thus, a *mode* of *robert the robot* would be treated as *mode r*, and a *mode* of *anything else* would be  
28271 treated as *mode w*.

28272 If the application calls *waitpid()* with a *pid* argument greater than zero, and it still has a  
28273 *popen()*ed stream open, it must ensure that *pid* does not refer to the process started by *popen()*.

28274 **FUTURE DIRECTIONS**

28275 None.

28276 **SEE ALSO**

28277 *pclose()*, *pipe()*, *sysconf()*, *system()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
28278 `<stdio.h>`, the Shell and Utilities volume of IEEE Std. 1003.1-200x

28279 **CHANGE HISTORY**

28280 First released in Issue 1. Derived from Issue 1 of the SVID.

28281 **Issue 4**

28282 The APPLICATION USAGE section is extended. Only notes about buffer flushing are retained  
28283 from Issue 3.

28284 The following changes are incorporated for alignment with the ISO POSIX-2 standard:

- 28285 • The function is no longer marked as an extension.
- 28286 • The type of arguments *command* and *mode* are changed from **char\*** to **const char\***.
- 28287 • The DESCRIPTION is completely rewritten for alignment with the ISO POSIX-2 standard,  
28288 although it describes essentially the same functionality as Issue 3.
- 28289 • The *sh* utility defined in the Shell and Utilities volume of IEEE Std. 1003.1-200x is no longer  
28290 required in all circumstances.
- 28291 • The ERRORS section is added.

28292 **Issue 5**

28293 A statement is added to the DESCRIPTION indicating that pipe streams are byte-oriented.

28294 **Issue 6**

28295 The following new requirements on POSIX implementations derive from alignment with the  
28296 Single UNIX Specification:

- 28297 • The optional [EMFILE] error condition is added.

28298 The following changes were made to align with the IEEE P1003.1a draft standard:

- 28299 • The DESCRIPTION is adjusted to reflect the behavior on systems that do not support the  
28300 Shell option.

28301 **NAME**

28302        **posix\_fadvise** — file advisory information

28303 **SYNOPSIS**

28304 ADV     #include <fcntl.h>

28305        int posix\_fadvise(int *fd*, off\_t *offset*, size\_t *len*, int *advice*);

28306

28307 **DESCRIPTION**

28308        The *posix\_fadvise()* function provides advice to the implementation on the expected behavior of  
 28309        the application with respect to the data in the file associated with the open file descriptor, *fd*,  
 28310        starting at *offset* and continuing for *len* bytes. The specified range need not currently exist in the  
 28311        file. If *len* is zero, all data following *offset* is specified. The implementation may use this  
 28312        information to optimize handling of the specified data. The *posix\_fadvise()* function has no effect  
 28313        on the semantics of other operations on the specified data, although it may affect the  
 28314        performance of other operations.

28315        The advice to be applied to the data is specified by the *advice* parameter and may be one of the  
 28316        following values:

28317        **POSIX\_FADV\_NORMAL**

28318            Specifies that the application has no advice to give on its behavior with respect to the  
 28319            specified data. It is the default characteristic if no advice is given for an open file.

28320        **POSIX\_FADV\_SEQUENTIAL**

28321            Specifies that the application expects to access the specified data sequentially from lower  
 28322            offsets to higher offsets.

28323        **POSIX\_FADV\_RANDOM**

28324            Specifies that the application expects to access the specified data in a random order.

28325        **POSIX\_FADV\_WILLNEED**

28326            Specifies that the application expects to access the specified data in the near future.

28327        **POSIX\_FADV\_DONTNEED**

28328            Specifies that the application expects that it will not access the specified data in the near  
 28329            future.

28330        **POSIX\_FADV\_NOREUSE**

28331            Specifies that the application expects to access the specified data once and then not reuse it  
 28332            thereafter.

28333        These values shall be defined in <fcntl.h>.

28334 **RETURN VALUE**

28335        Upon successful completion, *posix\_fadvise()* shall return zero; otherwise, an error number shall  
 28336        be returned to indicate the error.

28337 **ERRORS**

28338        The *posix\_fadvise()* function shall fail if:

28339        [EBADF]        The *fd* argument is not a valid file descriptor.

28340        [EINVAL]       The value of *advice* is invalid.

28341        [ESPIPE]       The *fd* argument is associated with a pipe or FIFO.

28342 **EXAMPLES**

28343           None.

28344 **APPLICATION USAGE**

28345           The *posix\_fadvise()* function is part of the Advisory Information option and need not be provided |  
28346           on all implementations.

28347 **RATIONALE**

28348           None.

28349 **FUTURE DIRECTIONS**

28350           None.

28351 **SEE ALSO**28352           *posix\_madvise()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <fcntl.h> |28353 **CHANGE HISTORY**

28354           First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

28355           In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

28356 **NAME**

28357 posix\_fallocate — file space control

28358 **SYNOPSIS**

28359 ADV #include <fcntl.h>

28360 int posix\_fallocate(int fd, off\_t offset, size\_t len);

28361

28362 **DESCRIPTION**

28363 The *posix\_fallocate()* function ensures that any required storage for regular file data starting at  
 28364 *offset* and continuing for *len* bytes is allocated on the file system storage media. If *posix\_fallocate()*  
 28365 returns successfully, subsequent writes to the specified file data shall not fail due to the lack of  
 28366 free space on the file system storage media.

28367 If the *offset+len* is beyond the current file size, then *posix\_fallocate()* shall adjust the file size to  
 28368 *offset+len*. Otherwise, the file size shall not be changed.

28369 It is implementation-defined whether a previous *posix\_fadvise()* call influences allocation  
 28370 strategy.

28371 Space allocated via *posix\_fallocate()* shall be freed by a successful call to *creat()* or *open()* that  
 28372 truncates the size of the file. Space allocated via *posix\_fallocate()* may be freed by a successful call  
 28373 to *ftruncate()* that reduces the file size to a size smaller than *offset+len*.

28374 **RETURN VALUE**

28375 Upon successful completion, *posix\_fallocate()* shall return zero; otherwise, an error number shall  
 28376 be returned to indicate the error.

28377 **ERRORS**

28378 The *posix\_fallocate()* function shall fail if:

- 28379 [EBADF] The *fd* argument is not a valid file descriptor.
- 28380 [EBADF] The *fd* argument references a file that was opened without write permission.
- 28381 [EFBIG] The value of *offset+len* is greater than the maximum file size.
- 28382 [EINTR] A signal was caught during execution.
- 28383 [EINVAL] The *len* argument was zero or the *offset* argument was less than zero.
- 28384 [EIO] An I/O error occurred while reading from or writing to a file system.
- 28385 [ENODEV] The *fd* argument does not refer to a regular file.
- 28386 [ENOSPC] There is insufficient free space remaining on the file system storage media.
- 28387 [ESPIPE] The *fd* argument is associated with a pipe or FIFO.

28388 **EXAMPLES**

28389 None.

28390 **APPLICATION USAGE**

28391 The *posix\_fallocate()* function is part of the Advisory Information option and need not be  
 28392 provided on all implementations.

28393 **RATIONALE**

28394 None.

28395 **FUTURE DIRECTIONS**

28396       None.

28397 **SEE ALSO**28398       *creat()*, *ftruncate()*, *open()*, *unlink()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
28399       <fcntl.h>28400 **CHANGE HISTORY**

28401       First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

28402       In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

28403 **NAME**

28404 posix\_madvise — memory advisory information and alignment control

28405 **SYNOPSIS**

28406 ADV #include <sys/mman.h>

28407 int posix\_madvise(void \*addr, size\_t len, int advice);

28408

28409 **DESCRIPTION**

28410 MF|SHM The *posix\_madvise()* function need only be supported if either the Memory Mapped Files or the  
28411 Shared Memory Objects options are supported.

28412 The *posix\_madvise()* function provides advice to the implementation on the expected behavior of  
28413 the application with respect to the data in the memory starting at address *addr*, and continuing  
28414 for *len* bytes. The implementation may use this information to optimize handling of the specified  
28415 data. The *posix\_madvise()* function has no effect on the semantics of access to memory in the  
28416 specified range, although it may affect the performance of access.

28417 The implementation may require that *addr* be a multiple of the page size, which is the value  
28418 returned by *sysconf()* when the name value *\_SC\_PAGESIZE* is used.

28419 The advice to be applied to the memory range is specified by the *advice* parameter and may be  
28420 one of the following values:

28421 POSIX\_MADV\_NORMAL

28422 Specifies that the application has no advice to give on its behavior with respect to the  
28423 specified range. It is the default characteristic if no advice is given for a range of memory.

28424 POSIX\_MADV\_SEQUENTIAL

28425 Specifies that the application expects to access the specified range sequentially from lower  
28426 addresses to higher addresses.

28427 POSIX\_MADV\_RANDOM

28428 Specifies that the application expects to access the specified range in a random order.

28429 POSIX\_MADV\_WILLNEED

28430 Specifies that the application expects to access the specified range in the near future.

28431 POSIX\_MADV\_DONTNEED

28432 Specifies that the application expects that it will not access the specified range in the near  
28433 future.

28434 These values shall be defined in <sys/mman.h>.

28435 **RETURN VALUE**

28436 Upon successful completion, *posix\_madvise()* shall return zero; otherwise, an error number shall  
28437 be returned to indicate the error.

28438 **ERRORS**

28439 The *posix\_madvise()* function shall fail if:

28440 [EINVAL] The value of *advice* is invalid.

28441 [ENOMEM] Addresses in the range starting at *addr* and continuing for *len* bytes are partly  
28442 or completely outside the range allowed for the address space of the calling  
28443 process.

28444 The *posix\_madvise()* function may fail if:

28445 [EINVAL] The value of *addr* is not a multiple of the value returned by *sysconf()* when the  
28446 name value *\_SC\_PAGESIZE* is used.

28447 [EINVAL] The value of *len* is zero.

28448 **EXAMPLES**

28449 None.

28450 **APPLICATION USAGE**

28451 The *posix\_madvise()* function is part of the Advisory Information option and need not be |  
28452 provided on all implementations.

28453 **RATIONALE**

28454 None.

28455 **FUTURE DIRECTIONS**

28456 None.

28457 **SEE ALSO**

28458 *mmap()*, *posix\_fadvise()*, *sysconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
28459 <**sys/mmap.h**>

28460 **CHANGE HISTORY**

28461 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

28462 IEEE PASC Interpretation 1003.1 #102 is included. |

28463 **NAME**

28464 posix\_mem\_offset — find offset and length of a mapped typed memory block

28465 **SYNOPSIS**

28466 TYM #include <sys/mman.h>

```
28467 int posix_mem_offset(const void *restrict addr, size_t len,
28468 off_t *restrict off, size_t *restrict contig_len,
28469 int *restrict fildes);
28470
```

28471 **DESCRIPTION**

28472 The *posix\_mem\_offset()* function returns in the variable pointed to by *off* a value that identifies  
 28473 the offset (or location), within a memory object, of the memory block currently mapped at *addr*.  
 28474 The function shall return in the variable pointed to by *fildes*, the descriptor used (via *mmap()*) to  
 28475 establish the mapping which contains *addr*. If that descriptor was closed since the mapping was  
 28476 established, the returned value of *fildes* shall be *-1*. The *len* argument specifies the length of the  
 28477 block of the memory object the user wishes the offset for; upon return, the value pointed to by  
 28478 *contig\_len* shall equal either *len*, or the length of the largest contiguous block of the memory  
 28479 object that is currently mapped to the calling process starting at *addr*, whichever is smaller.

28480 If the memory object mapped at *addr* is a typed memory object, then if the *off* and *contig\_len*  
 28481 values obtained by calling *posix\_mem\_offset()* are used in a call to *mmap()* with a file descriptor  
 28482 that refers to the same memory pool as *fildes* (either through the same port or through a different  
 28483 port), and that was opened with neither the *POSIX\_TYPED\_MEM\_ALLOCATE* nor the  
 28484 *POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG* flag, the typed memory area that is mapped shall  
 28485 be exactly the same area that was mapped at *addr* in the address space of the process that called  
 28486 *posix\_mem\_offset()*.

28487 If the memory object specified by *fildes* is not a typed memory object, then the behavior of this  
 28488 function is implementation-defined.

28489 **RETURN VALUE**

28490 Upon successful completion, the *posix\_mem\_offset()* function shall return zero; otherwise, the  
 28491 corresponding error status value shall be returned.

28492 **ERRORS**

28493 The *posix\_mem\_offset()* function shall fail if:

28494 [EACCES] The process has not mapped a memory object supported by this function at  
 28495 the given address *addr*.

28496 This function shall not return an error code of [EINTR].

28497 **EXAMPLES**

28498 None.

28499 **APPLICATION USAGE**

28500 None.

28501 **RATIONALE**

28502 None.

28503 **FUTURE DIRECTIONS**

28504 None.

28505 **SEE ALSO**

28506 *mmap()*, *posix\_typed\_mem\_open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
28507 `<sys/mman.h>`

28508 **CHANGE HISTORY**

28509 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

28510 **NAME**

28511 posix\_memalign — aligned memory allocation

28512 **SYNOPSIS**

28513 ADV #include <stdlib.h>

28514 int posix\_memalign(void \*\*memptr, size\_t alignment, size\_t size);

28515

28516 **DESCRIPTION**

28517 The *posix\_memalign()* function allocates *size* bytes aligned on a boundary specified by *alignment*,  
 28518 and returns a pointer to the allocated memory in *memptr*. The value of *alignment* shall be a  
 28519 multiple of *sizeof(void\*)*, that is also a power of two. Upon successful completion, the value  
 28520 pointed to by *memptr* shall be a multiple of *alignment*.

28521 CX The *free()* function shall deallocate memory that has previously been allocated by  
 28522 *posix\_memalign()*.

28523 **RETURN VALUE**

28524 Upon successful completion, *posix\_memalign()* shall return zero; otherwise, an error number  
 28525 shall be returned to indicate the error.

28526 **ERRORS**

28527 The *posix\_memalign()* function shall fail if:

28528 [EINVAL] The value of the alignment parameter is not a power of two multiple of  
 28529 *sizeof(void\*)*.

28530 [ENOMEM] There is insufficient memory available with the requested alignment.

28531 **EXAMPLES**

28532 None.

28533 **APPLICATION USAGE**

28534 The *posix\_memalign()* function is part of the Advisory Information option and need not be  
 28535 provided on all implementations.

28536 **RATIONALE**

28537 None.

28538 **FUTURE DIRECTIONS**

28539 None.

28540 **SEE ALSO**

28541 *free()*, *malloc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

28542 **CHANGE HISTORY**

28543 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

28544 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

28545 **NAME**28546 posix\_spawn, posix\_spawnnp — spawn a process (**REALTIME**)28547 **SYNOPSIS**

```
28548 SPN #include <spawn.h>
28549
28549 int posix_spawn(pid_t *restrict pid, const char *restrict path,
28550 const posix_spawn_file_actions_t *file_actions,
28551 const posix_spawnattr_t *restrict attrp,
28552 char *restrict const argv[restrict],
28553 char *restrict const envp[restrict]);
28554 int posix_spawnnp(pid_t *restrict pid, const char *restrict file,
28555 const posix_spawn_file_actions_t *file_actions,
28556 const posix_spawnattr_t *restrict attrp, char *const argv[restrict],
28557 char * const envp[restrict]);
28558
```

28559 **DESCRIPTION**

28560 The *posix\_spawn()* and *posix\_spawnnp()* functions shall create a new process (child process) from  
 28561 the specified process image. The new process image is constructed from a regular executable file  
 28562 called the new process image file.

28563 When a C program is executed as the result of this call, it shall be entered as a C language  
 28564 function call as follows:

```
28565 int main(int argc, char *argv[]);
```

28566 where *argc* is the argument count and *argv* is an array of character pointers to the arguments  
 28567 themselves. In addition, the following variable:

```
28568 extern char **environ;
```

28569 is initialized as a pointer to an array of character pointers to the environment strings.

28570 The argument *argv* is an array of character pointers to null-terminated strings. The last member  
 28571 of this array shall be a null pointer and is not counted in *argc*. These strings constitute the  
 28572 argument list available to the new process image. The value in *argv[0]* should point to a file  
 28573 name that is associated with the process image being started by the *posix\_spawn()* or  
 28574 *posix\_spawnnp()* function.

28575 The argument *envp* is an array of character pointers to null-terminated strings. These strings  
 28576 constitute the environment for the new process image. The environment array is terminated by a  
 28577 null pointer.

28578 The number of bytes available for the child process' combined argument and environment lists  
 28579 is {ARG\_MAX}. The implementation shall specify in the system documentation (see the Base  
 28580 Definitions volume of IEEE Std. 1003.1-200x, Chapter 2, Conformance) whether any list  
 28581 overhead, such as length words, null terminators, pointers, or alignment bytes, is included in  
 28582 this total.

28583 The *path* argument to *posix\_spawn()* is a path name that identifies the new process image file to  
 28584 execute.

28585 The *file* parameter to *posix\_spawnnp()* shall be used to construct a path name that identifies the  
 28586 new process image file. If the *file* parameter contains a slash character, the *file* parameter shall be  
 28587 used as the path name for the new process image file. Otherwise, the path prefix for this file shall  
 28588 be obtained by a search of the directories passed as the environment variable *PATH* (see the Base  
 28589 Definitions volume of IEEE Std. 1003.1-200x, Chapter 8, Environment Variables). If this  
 28590 environment variable is not defined, the results of the search are implementation-defined.

28591 If *file\_actions* is a null pointer, then file descriptors open in the calling process shall remain open  
 28592 in the child process, except for those whose close-on-exec flag FD\_CLOEXEC is set (see *fcntl()*).  
 28593 For those file descriptors that remain open, all attributes of the corresponding open file  
 28594 descriptions, including file locks (see *fcntl()*), shall remain unchanged.

28595 If *file\_actions* is not NULL, then the file descriptors open in the child process shall be those open  
 28596 in the calling process as modified by the spawn file actions object pointed to by *file\_actions* and  
 28597 the FD\_CLOEXEC flag of each remaining open file descriptor after the spawn file actions have  
 28598 been processed. The effective order of processing the spawn file actions shall be:

- 28599 1. The set of open file descriptors for the child process shall initially be the same set as is  
 28600 open for the calling process. All attributes of the corresponding open file descriptions,  
 28601 including file locks (see *fcntl()*), shall remain unchanged.
- 28602 2. The signal mask, signal default actions, and the effective user and group IDs for the child  
 28603 process shall be changed as specified in the attributes object referenced by *attrp*.
- 28604 3. The file actions specified by the spawn file actions object shall be performed in the order in  
 28605 which they were added to the spawn file actions object.
- 28606 4. Any file descriptor that has its FD\_CLOEXEC flag set (see *fcntl()*) shall be closed.

28607 The **posix\_spawnattr\_t** spawn attributes object type is defined in **<spawn.h>**. It shall contain at  
 28608 least the attributes defined below.

28609 If the POSIX\_SPAWN\_SETPGROUP flag is set in the *spawn\_flags* attribute of the object  
 28610 referenced by *attrp*, and the *spawn-pgroup* attribute of the same object is non-zero, then the  
 28611 child's process group shall be as specified in the *spawn-pgroup* attribute of the object referenced  
 28612 by *attrp*.

28613 As a special case, if the POSIX\_SPAWN\_SETPGROUP flag is set in the *spawn\_flags* attribute of  
 28614 the object referenced by *attrp*, and the *spawn-pgroup* attribute of the same object is set to zero,  
 28615 then the child shall be in a new process group with a process group ID equal to its process ID.

28616 If the POSIX\_SPAWN\_SETPGROUP flag is not set in the *spawn\_flags* attribute of the object  
 28617 referenced by *attrp*, the new child process shall inherit the parent's process group.

28618 PS If the POSIX\_SPAWN\_SETSCHEDPARAM flag is set in the *spawn\_flags* attribute of the object  
 28619 referenced by *attrp*, but POSIX\_SPAWN\_SETSCHEDULER is not set, the new process image  
 28620 shall initially have the scheduling policy of the calling process with the scheduling parameters  
 28621 specified in the *spawn-schedparam* attribute of the object referenced by *attrp*.

28622 If the POSIX\_SPAWN\_SETSCHEDULER flag is set in *spawn\_flags* attribute of the object  
 28623 referenced by *attrp* (regardless of the setting of the POSIX\_SPAWN\_SETSCHEDPARAM flag),  
 28624 the new process image shall initially have the scheduling policy specified in the *spawn-*  
 28625 *schedpolicy* attribute of the object referenced by *attrp* and the scheduling parameters specified in  
 28626 the *spawn-schedparam* attribute of the same object.

28627 The POSIX\_SPAWN\_RESETEUIDS flag in the *spawn\_flags* attribute of the object referenced by *attrp*  
 28628 governs the effective user ID of the child process. If this flag is not set, the child process inherits  
 28629 the parent process' effective user ID. If this flag is set, the child process' effective user ID is reset  
 28630 to the parent's real user ID. In either case, if the set-user-ID mode bit of the new process image  
 28631 file is set, the effective user ID of the child process will become that file's owner ID before the  
 28632 new process image begins execution.

28633 The POSIX\_SPAWN\_RESETEGID flag in the *spawn\_flags* attribute of the object referenced by *attrp*  
 28634 also governs the effective group ID of the child process. If this flag is not set, the child process  
 28635 inherits the parent process' effective group ID. If this flag is set, the child process' effective group  
 28636 ID is reset to the parent's real group ID. In either case, if the set-group-ID mode bit of the new

28637 process image file is set, the effective group ID of the child process will become that file's group  
28638 ID before the new process image begins execution.

28639 If the POSIX\_SPAWN\_SETSIGMASK flag is set in the *spawn-flags* attribute of the object  
28640 referenced by *attrp*, the child process shall initially have the signal mask specified in the *spawn-*  
28641 *sigmask* attribute of the object referenced by *attrp*.

28642 If the POSIX\_SPAWN\_SETSIGDEF flag is set in the *spawn-flags* attribute of the object referenced  
28643 by *attrp*, the signals specified in the *spawn-sigdefault* attribute of the same object shall be set to  
28644 their default actions in the child process. Signals set to the default action in the parent process  
28645 shall be set to the default action in the child process.

28646 Signals set to be caught by the calling process shall be set to the default action in the child  
28647 process.

28648 Signals set to be ignored by the calling process image shall be set to be ignored by the child  
28649 process, unless otherwise specified by the POSIX\_SPAWN\_SETSIGDEF flag being set in the  
28650 *spawn-flags* attribute of the object referenced by *attrp* and the signals being indicated in the  
28651 *spawn-sigdefault* attribute of the object referenced by *attrp*.

28652 If the value of the *attrp* pointer is NULL, then the default values are used.

28653 All process attributes, other than those influenced by the attributes set in the object referenced  
28654 by *attrp* as specified above or by the file descriptor manipulations specified in *file\_actions*, shall  
28655 appear in the new process image as though *fork()* had been called to create a child process and  
28656 then a member of the *exec* family of functions had been called by the child process to execute the  
28657 new process image.

28658 THR It is implementation-defined whether the fork handlers are run when *posix\_spawn()* or  
28659 *posix\_spawnp()* is called.

#### 28660 RETURN VALUE

28661 Upon successful completion, *posix\_spawn()* and *posix\_spawnp()* shall return the process ID of the  
28662 child process to the parent process, in the variable pointed to by a non-NULL *pid* argument, and  
28663 shall return zero as the function return value. Otherwise, no child process shall be created, the  
28664 value stored into the variable pointed to by a non-NULL *pid* is unspecified, and an error number  
28665 shall be returned as the function return value to indicate the error. If the *pid* argument is a null  
28666 pointer, the process ID of the child is not returned to the caller.

#### 28667 ERRORS

28668 The *posix\_spawn()* and *posix\_spawnp()* functions may fail if:

28669 [EINVAL] The value specified by *file\_actions* or *attrp* is invalid.

28670 If this error occurs after the calling process successfully returns from the *posix\_spawn()* or  
28671 *posix\_spawnp()* function, the child process may exit with exit status 127.

28672 If *posix\_spawn()* or *posix\_spawnp()* fail for any of the reasons that would cause *fork()* or one of  
28673 the *exec* family of functions to fail, an error value shall be returned as described by *fork()* and  
28674 *exec*, respectively (or, if the error occurs after the calling process successfully returns, the child  
28675 process exits with exit status 127).

28676 If POSIX\_SPAWN\_SETPGROUP is set in the *spawn-flags* attribute of the object referenced by  
28677 *attrp*, and *posix\_spawn()* or *posix\_spawnp()* fails while changing the child's process group, an  
28678 error value shall be returned as described by *setpgid()* (or, if the error occurs after the calling  
28679 process successfully returns, the child process exits with exit status 127).

28680 PS If POSIX\_SPAWN\_SETSCHEDPARAM is set and POSIX\_SPAWN\_SETSCHEDULER is not set  
28681 in the *spawn-flags* attribute of the object referenced by *attrp*, then if *posix\_spawn()* or

28682 *posix\_spawnnp()* fails for any of the reasons that would cause *sched\_setparam()* to fail, an error  
 28683 value shall be returned as described by *sched\_setparam()* (or, if the error occurs after the calling  
 28684 process successfully returns, the child process exits with exit status 127).

28685 If POSIX\_SPAWN\_SETSCHEDULER is set in the *spawn-flags* attribute of the object referenced by  
 28686 *attrp*, and if *posix\_spawn()* or *posix\_spawnnp()* fails for any of the reasons that would cause  
 28687 *sched\_setscheduler()* to fail, an error value shall be returned as described by *sched\_setscheduler()*  
 28688 (or, if the error occurs after the calling process successfully returns, the child process exits with  
 28689 exit status 127)>

28690 If the *file\_actions* argument is not NULL, and specifies any *close*, *dup2*, or *open* actions to be  
 28691 performed, and if *posix\_spawn()* or *posix\_spawnnp()* fails for any of the reasons that would cause  
 28692 *close()*, *dup2()*, or *open()* to fail, an error value shall be returned as described by *close()*, *dup2()*,  
 28693 and *open()*, respectively (or, if the error occurs after the calling process successfully returns, the  
 28694 child process exits with exit status 127). An open file action may, by itself, result in any of the  
 28695 errors described by *close()* or *dup2()*, in addition to those described by *open()*.

28696 **EXAMPLES**

28697 None.

28698 **APPLICATION USAGE**

28699 These functions are part of the Spawn option and need not be provided on all implementations.

28700 **RATIONALE**

28701 The POSIX *fork()* function is difficult or impossible to implement without swapping or dynamic  
 28702 address translation for the following reasons:

- 28703 • Swapping is generally too slow for a realtime environment.
- 28704 • Dynamic address translation is not available everywhere POSIX might be useful.
- 28705 • Processes are too useful to simply option out of POSIX whenever it must run without  
 28706 address translation or other MMU services,

28707 POSIX needs process creation and file execution primitives that can be efficiently implemented  
 28708 without address translation or other MMU services.

28709 This function shall be called *posix\_spawn()*. A closely related function, *posix\_spawnnp()*, is  
 28710 included for completeness.

28711 The *posix\_spawn()* function is implementable as a library routine, but both *posix\_spawn()* and  
 28712 *posix\_spawnnp()* are designed as kernel operations. Also, although they may be an efficient  
 28713 replacement for many *fork()/exec* pairs, their goal is to provide useful process creation  
 28714 primitives for systems that have difficulty with *fork()*, not to provide drop-in replacements for  
 28715 *fork()/exec*.

28716 This view of the role of *posix\_spawn()* and *posix\_spawnnp()* influenced the design of their API. It  
 28717 does not attempt to provide the full functionality of *fork()/exec* in which arbitrary user-specified  
 28718 operations of any sort are permitted between the creation of the child process and the execution  
 28719 of the new process image; any attempt to reach that level would need to provide a programming  
 28720 language as parameters. Instead, *posix\_spawn()* and *posix\_spawnnp()* are process creation  
 28721 primitives like the *Start\_Process* and *Start\_Process\_Search* Ada language bindings package  
 28722 *POSIX\_Process\_Primitives* and also like those in many operating systems that are not UNIX  
 28723 systems, but with some POSIX-specific additions.

28724 To achieve its coverage goals, *posix\_spawn()* and *posix\_spawnnp()* have control of six types of  
 28725 inheritance: file descriptors, process group ID, user and group ID, signal mask, scheduling, and  
 28726 whether each signal ignored in the parent will remain ignored in the child, or be reset to its  
 28727 default action in the child.

28728 Control of file descriptors is required to allow an independently written child process image to  
 28729 access data streams opened by and even generated or read by the parent process without being  
 28730 specifically coded to know which parent files and file descriptors are to be used. Control of the  
 28731 process group ID is required to control how the child process' job control relates to that of the  
 28732 parent.

28733 Control of the signal mask and signal defaulting is sufficient to support the implementation of  
 28734 *system()*. Although support for *system()* is not explicitly one of the goals for *posix\_spawn()* and  
 28735 *posix\_spawnnp()*, it is covered under the "at least 50%" coverage goal.

28736 The intention is that the normal file descriptor inheritance across *fork()*, the subsequent effect of  
 28737 the specified spawn file actions, and the normal file descriptor inheritance across one of the *exec*  
 28738 family of functions should fully specify open file inheritance. The implementation need make no  
 28739 decisions regarding the set of open file descriptors when the child process image begins  
 28740 execution, those decisions having already been made by the caller and expressed as the set of  
 28741 open file descriptors and their *FD\_CLOEXEC* flags at the time of the call and the spawn file  
 28742 actions object specified in the call. We have been assured that in cases where the POSIX  
 28743 *Start\_Process* Ada primitives have been implemented in a library, this method of controlling file  
 28744 descriptor inheritance may be implemented very easily.

28745 We can identify several problems with *posix\_spawn()* and *posix\_spawnnp()*, but there does not  
 28746 appear to be a solution that introduces fewer problems. Environment modification for child  
 28747 process attributes not specifiable via the *attrp* or *file\_actions* arguments must be done in the  
 28748 parent process, and since the parent generally wants to save its context, it is more costly than  
 28749 similar functionality with *fork()/exec*. It is also complicated to modify the environment of a  
 28750 multi-threaded process temporarily, since all threads must agree when it is safe for the  
 28751 environment to be changed. However, this cost is only borne by those invocations of  
 28752 *posix\_spawn()* and *posix\_spawnnp()* that use the additional functionality. Since extensive  
 28753 modifications are not the usual case, and are particularly unlikely in time-critical code, keeping  
 28754 much of the environment control out of *posix\_spawn()* and *posix\_spawnnp()* is appropriate design.

28755 The *posix\_spawn()* and *posix\_spawnnp()* functions do not have all the power of *fork()/exec*. This is  
 28756 to be expected. The *fork()* function is a wonderfully powerful operation. We do not expect to  
 28757 duplicate its functionality in a simple, fast function with no special hardware requirements. It is  
 28758 worth noting that *posix\_spawn()* and *posix\_spawnnp()* are very similar to the process creation  
 28759 operations on many operating systems that are not UNIX systems.

## 28760 Requirements

28761 The requirements for *posix\_spawn()* and *posix\_spawnnp()* are:

- 28762 • They must be implementable without an MMU or unusual hardware.
- 28763 • They must be compatible with existing POSIX standards.

28764 Additional goals are:

- 28765 • They should be efficiently implementable.
- 28766 • They should be able to replace at least 50% of typical executions of *fork()*.
- 28767 • A system with *posix\_spawn()* and *posix\_spawnnp()* and without *fork()* should be useful, at least  
 28768 for realtime applications.
- 28769 • A system with *fork()* and the *exec* family should be able to implement *posix\_spawn()* and  
 28770 *posix\_spawnnp()* as library routines.

28771 **Two-Syntax**

28772 POSIX *exec* has several calling sequences with approximately the same functionality. These  
 28773 appear to be required for compatibility with existing practice. Since the existing practice for the  
 28774 *posix\_spawn\*()* functions is otherwise substantially unlike POSIX, we feel that simplicity  
 28775 outweighs compatibility. There are, therefore, only two names for the *posix\_spawn\*()* functions.

28776 The parameter list does not differ between *posix\_spawn()* and *posix\_spawnp()*; *posix\_spawnp()*  
 28777 interprets the second parameter more elaborately than *posix\_spawn()*.

28778 **Compatibility with POSIX.5 (Ada)**

28779 The *Start\_Process* and *Start\_Process\_Search* procedures from the *POSIX\_Process\_Primitives*  
 28780 package from the Ada language binding to POSIX.1 encapsulate *fork()* and *exec* functionality in a  
 28781 manner similar to that of *posix\_spawn()* and *posix\_spawnp()*. Originally, in keeping with our  
 28782 simplicity goal, the standard developers had limited the capabilities of *posix\_spawn()* and  
 28783 *posix\_spawnp()* to a subset of the capabilities of *Start\_Process* and *Start\_Process\_Search*; certain  
 28784 non-default capabilities were not supported. However, based on suggestions by the ballot group  
 28785 to improve file descriptor mapping or drop it, and on the advice of an Ada Language Bindings  
 28786 working group member, the standard developers decided that *posix\_spawn()* and *posix\_spawnp()*  
 28787 should be sufficiently powerful to implement *Start\_Process* and *Start\_Process\_Search*. The  
 28788 rationale is that if the Ada language binding to such a primitive had already been approved as  
 28789 an IEEE standard, there can be little justification for not approving the functionally-equivalent  
 28790 parts of a C binding. The only three capabilities provided by *posix\_spawn()* and *posix\_spawnp()*  
 28791 that are not provided by *Start\_Process* and *Start\_Process\_Search* are optionally specifying the  
 28792 child's process group ID, the set of signals to be reset to default signal handling in the child  
 28793 process, and the child's scheduling policy and parameters.

28794 For the Ada language binding for *Start\_Process* to be implemented with *posix\_spawn()*, that  
 28795 binding would need to explicitly pass an empty signal mask and the parent's environment to  
 28796 *posix\_spawn()* whenever the caller of *Start\_Process* allowed these arguments to default, since  
 28797 *posix\_spawn()* does not provide such defaults. The ability of *Start\_Process* to mask user-specified  
 28798 signals during its execution is functionally unique to the Ada language binding and must be  
 28799 dealt with in the binding separately from the call to *posix\_spawn()*.

28800 **Process Group**

28801 The process group inheritance field can be used to join the child process with an existing process  
 28802 group. By assigning a value of zero to the *spawn-pgroup* attribute of the object referenced by  
 28803 *attrp*, the *setpgid()* mechanism will place the child process in a new process group.

28804 **Threads**

28805 Without the *posix\_spawn()* and *posix\_spawnp()* functions, systems without address translation  
 28806 can still use threads to give an abstraction of concurrency. In many cases, thread creation  
 28807 suffices, but it is not always a good substitute. The *posix\_spawn()* and *posix\_spawnp()* functions  
 28808 are considerably "heavier" than thread creation. Processes have several important attributes that  
 28809 threads do not. Even without address translation, a process may have base-and-bound memory  
 28810 protection. Each process has a process environment including security attributes and file  
 28811 capabilities, and powerful scheduling attributes. Processes abstract the behavior of non-  
 28812 uniform-memory-architecture multi-processors better than threads, and they are more  
 28813 convenient to use for activities that are not closely linked.

28814 The *posix\_spawn()* and *posix\_spawnp()* functions may not bring support for multiple processes to  
 28815 every configuration. Process creation is not the only piece of operating system support required  
 28816 to support multiple processes. The total cost of support for multiple processes may be quite high

28817 in some circumstances. Existing practice shows that support for multiple processes is  
 28818 uncommon and threads are common among “tiny kernels”. There should, therefore, probably  
 28819 continue to be AEPs for operating systems with only one process.

#### 28820 **Asynchronous Error Notification**

28821 A library implementation of *posix\_spawn()* or *posix\_spawnp()* may not be able to detect all  
 28822 possible errors before it forks the child process. IEEE Std. 1003.1-200x provides for an error  
 28823 indication returned from a child process which could not successfully complete the spawn  
 28824 operation via a special exit status which may be detected using the status value returned by  
 28825 *wait()* and *waitpid()*.

28826 The *stat\_val* interface and the macros used to interpret it are not well suited to the purpose of  
 28827 returning API errors, but they are the only path available to a library implementation. Thus, an  
 28828 implementation may cause the child process to exit with exit status 127 for any error detected  
 28829 during the spawn process after the *posix\_spawn()* or *posix\_spawnp()* function has successfully  
 28830 returned.

28831 The standard developers had proposed using two additional macros to interpret *stat\_val*. The  
 28832 first, WIFSPAWNFAIL, would have detected a status that indicated that the child exited because  
 28833 of an error detected during the *posix\_spawn()* or *posix\_spawnp()* operations rather than during  
 28834 actual execution of the child process image; the second, WSPAWNERRNO, would have  
 28835 extracted the error value if WIFSPAWNFAIL indicated a failure. Unfortunately, the ballot group  
 28836 strongly opposed this because it would make a library implementation of *posix\_spawn()* or  
 28837 *posix\_spawnp()* dependent on kernel modifications to *waitpid()* to be able to embed special  
 28838 information in *stat\_val* to indicate a spawn failure.

28839 The 8 bits of child process exit status that are guaranteed by IEEE Std. 1003.1-200x to be  
 28840 accessible to the waiting parent process are insufficient to disambiguate a spawn error from any  
 28841 other kind of error that may be returned by an arbitrary process image. No other bits of the exit  
 28842 status are required to be visible in *stat\_val*, so these macros could not be strictly implemented at  
 28843 the library level. Reserving an exit status of 127 for such spawn errors is consistent with the use  
 28844 of this value by *system()* and *popen()* to signal failures in these operations that occur after the  
 28845 function has returned but before a shell is able to execute. The exit status of 127 does not  
 28846 uniquely identify this class of error, nor does it provide any detailed information on the nature  
 28847 of the failure. Note that a kernel implementation of *posix\_spawn()* or *posix\_spawnp()* is permitted  
 28848 (and encouraged) to return any possible error as the function value, thus providing more  
 28849 detailed failure information to the parent process.

28850 Thus, no special macros are available to isolate asynchronous *posix\_spawn()* or *posix\_spawnp()*  
 28851 errors. Instead, errors detected by the *posix\_spawn()* or *posix\_spawnp()* operations in the context  
 28852 of the child process before the new process image executes are reported by setting the child’s  
 28853 exit status to 127. The calling process may use the WIFEXITED and WEXITSTATUS macros on  
 28854 the *stat\_val* stored by the *wait()* or *waitpid()* functions to detect spawn failures to the extent that  
 28855 other status values with which the child process image may exit (before the parent can  
 28856 conclusively determine that the child process image has begun execution) are distinct from exit  
 28857 status 127.

#### 28858 **FUTURE DIRECTIONS**

28859 None.

#### 28860 **SEE ALSO**

28861 *alarm()*, *chmod()*, *close()*, *dup()*, *exec*, *exit()*, *fcntl()*, *fork()*, *kill()*, *open()*,  
 28862 *posix\_spawn\_file\_actions\_addclose()*, *posix\_spawn\_file\_actions\_adddup2()*,  
 28863 *posix\_spawn\_file\_actions\_addopen()*, *posix\_spawn\_file\_actions\_destroy()*,  
 28864 *posix\_spawn\_file\_actions\_init()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*,

28865 *posix\_spawnattr\_getsigdefault()*, *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getpgroup()*,  
28866 *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*, *posix\_spawnattr\_getsigmask()*,  
28867 *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*, *posix\_spawnattr\_setpgroup()*,  
28868 *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
28869 *sched\_setparam()*, *sched\_setscheduler()*, *setpgid()*, *setuid()*, *stat()*, *times()*, *wait()*, the Base  
28870 Definitions volume of IEEE Std. 1003.1-200x, <spawn.h>

**28871 CHANGE HISTORY**

28872 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

28873 IEEE PASC Interpretation 1003.1 #103 is included, noting that the signal default actions are  
28874 changed as well as the signal mask in step 2.

28875 **NAME**

28876 posix\_spawn\_file\_actions\_addclose, posix\_spawn\_file\_actions\_addopen — add close or open  
 28877 action to spawn file actions object (**REALTIME**)

28878 **SYNOPSIS**

```
28879 SPN #include <spawn.h>

28880 int posix_spawn_file_actions_addclose(posix_spawn_file_actions_t *
28881 file_actions, int fildes);
28882 int posix_spawn_file_actions_addopen(posix_spawn_file_actions_t *restrict
28883 file_actions, int fildes, const char *restrict path,
28884 int oflag, mode_t mode);
28885
```

28886 **DESCRIPTION**

28887 A spawn file actions object is of type **posix\_spawn\_file\_actions\_t** (defined in **<spawn.h>**) and is  
 28888 used to specify a series of actions to be performed by a *posix\_spawn()* or *posix\_spawnp()*  
 28889 operation in order to arrive at the set of open file descriptors for the child process given the set of  
 28890 open file descriptors of the parent. IEEE Std. 1003.1-200x does not define comparison or  
 28891 assignment operators for the type **posix\_spawn\_file\_actions\_t**.

28892 A spawn file actions object, when passed to *posix\_spawn()* or *posix\_spawnp()*, shall specify how  
 28893 the set of open file descriptors in the calling process is transformed into a set of potentially open  
 28894 file descriptors for the spawned process. This transformation shall be as if the specified sequence  
 28895 of actions was performed exactly once, in the context of the spawned process (prior to execution  
 28896 of the new process image), in the order in which the actions were added to the object;  
 28897 additionally, when the new process image is executed, any file descriptor (from this new set)  
 28898 which has its **FD\_CLOEXEC** flag set will be closed (see *posix\_spawn()*).

28899 The *posix\_spawn\_file\_actions\_addclose()* function adds a *close* action to the object referenced by  
 28900 *file\_actions* that will cause the file descriptor *fildes* to be closed (as if *close(fildes)* had been called)  
 28901 when a new process is spawned using this file actions object.

28902 The *posix\_spawn\_file\_actions\_addopen()* function adds an *open* action to the object referenced by  
 28903 *file\_actions* that will cause the file named by *path* to be opened (as if *open(path, oflag, mode)* had  
 28904 been called, and the returned file descriptor, if not *fildes*, had been changed to *fildes*) when a new  
 28905 process is spawned using this file actions object. If *fildes* was already an open file descriptor, it  
 28906 shall be closed before the new file is opened.

28907 **Notes to Reviewers**

28908 *This section with side shading will not appear in the final copy. - Ed.*

28909 D3, XSH, ERN 448 says the description of the *posix\_spawn\_file\_actions\_addopen* function does  
 28910 not say whether the function has to make a copy of the path parameter or whether it can store  
 28911 the pointer and assume the application does not destroy the copy of the string. Add to the  
 28912 description: "The string pointed to by path can become invalid so the function has to make a  
 28913 copy."

28914 **RETURN VALUE**

28915 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
 28916 be returned to indicate the error.

28917 **ERRORS**

28918 These functions shall fail if:

28919 [EBADF] The value specified by *fildes* is negative or greater than or equal to  
 28920 {OPEN\_MAX}.

28921 These functions may fail if:

28922 [EINVAL] The value specified by *file\_actions* is invalid.

28923 [ENOMEM] Insufficient memory exists to add to the spawn file actions object.

28924 It shall not be considered an error for the *files* argument passed to these functions to specify a  
 28925 file descriptor for which the specified operation could not be performed at the time of the call.  
 28926 Any such error will be detected when the associated file actions object is later used during a  
 28927 *posix\_spawn()* or *posix\_spawnnp()* operation.

28928 **EXAMPLES**

28929 None.

28930 **APPLICATION USAGE**

28931 These functions are part of the Spawn option and need not be provided on all implementations.

28932 **RATIONALE**

28933 A spawn file actions object may be initialized to contain an ordered sequence of *close()*, *dup2()*,  
 28934 and *open()* operations to be used by *posix\_spawn()* or *posix\_spawnnp()* to arrive at the set of open  
 28935 file descriptors inherited by the spawned process from the set of open file descriptors in the  
 28936 parent at the time of the *posix\_spawn()* or *posix\_spawnnp()* call. It had been suggested that the  
 28937 *close()* and *dup2()* operations alone are sufficient to rearrange file descriptors, and that files  
 28938 which need to be opened for use by the spawned process can be handled either by having the  
 28939 calling process open them before the *posix\_spawn()* or *posix\_spawnnp()* call (and close them after),  
 28940 or by passing file names to the spawned process (in *argv*) so that it may open them itself. The  
 28941 standard developers recommend that applications use one of these two methods when practical,  
 28942 since detailed error status on a failed open operation is always available to the application this  
 28943 way. However, the standard developers feel that allowing a spawn file actions object to specify  
 28944 open operations is still appropriate because:

- 28945 1. It is consistent with equivalent POSIX.5 (Ada) functionality.
- 28946 2. It supports the I/O redirection paradigm commonly employed by POSIX programs  
 28947 designed to be invoked from a shell. When such a program is the child process, it may not  
 28948 be designed to open files on its own.
- 28949 3. It allows file opens that might otherwise fail or violate file ownership/access rights if  
 28950 executed by the parent process.

28951 Regarding 2. above, note that the spawn open file action provides to *posix\_spawn()* and  
 28952 *posix\_spawnnp()* the same capability that the shell redirection operators provide to *system()*, only  
 28953 without the intervening execution of a shell; for example:

```
28954 system ("myprog <file1 3<file2");
```

28955 Regarding 3. above, note that if the calling process needs to open one or more files for access by  
 28956 the spawned process, but has insufficient spare file descriptors, then the open action is necessary  
 28957 to allow the *open()* to occur in the context of the child process after other file descriptors have  
 28958 been closed (that must remain open in the parent).

28959 Additionally, if a parent is executed from a file having a “set-user-id” mode bit set and the  
 28960 POSIX\_SPAWN\_RESETEUIDS flag is set in the spawn attributes, a file created within the parent  
 28961 process will (possibly incorrectly) have the parent’s effective user ID as its owner, whereas a file  
 28962 created via an *open()* action during *posix\_spawn()* or *posix\_spawnnp()* will have the parent’s real  
 28963 ID as its owner; and an open by the parent process may successfully open a file to which the real  
 28964 user should not have access or fail to open a file to which the real user should have access.

28965 **File Descriptor Mapping**

28966 The standard developers had originally proposed using an array which specified the mapping of  
28967 child file descriptors back to those of the parent. It was pointed out by the ballot group that it is  
28968 not possible to reshuffle file descriptors arbitrarily in a library implementation of *posix\_spawn()*  
28969 or *posix\_spawnnp()* without provision for one or more spare file descriptor entries (which simply  
28970 may not be available). Such an array requires that an implementation develop a complex  
28971 strategy to achieve the desired mapping without inadvertently closing the wrong file descriptor  
28972 at the wrong time.

28973 It was noted by a member of the Ada Language Bindings working group that the approved Ada  
28974 Language *Start\_Process* family of POSIX process primitives use a caller-specified set of file  
28975 actions to alter the normal *fork()/exec* semantics for inheritance of file descriptors in a very  
28976 flexible way, yet no such problems exist because the burden of determining how to achieve the  
28977 final file descriptor mapping is completely on the application. Furthermore, although the file  
28978 actions interface appears frightening at first glance, it is actually quite simple to implement in  
28979 either a library or the kernel.

28980 **FUTURE DIRECTIONS**

28981 None.

28982 **SEE ALSO**

28983 *close()*, *dup()*, *open()*, *posix\_spawn()*, *posix\_spawn\_file\_actions\_adddup2()*,  
28984 *posix\_spawn\_file\_actions\_destroy()*, *posix\_spawnnp()*, the Base Definitions volume of  
28985 IEEE Std. 1003.1-200x, <spawn.h>

28986 **CHANGE HISTORY**

28987 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

28988 **NAME**

28989 posix\_spawn\_file\_actions\_adddup2 — add dup2 action to spawn file actions object  
 28990 (**REALTIME**)

28991 **SYNOPSIS**

```
28992 SPN #include <spawn.h>
28993
28994 int posix_spawn_file_actions_adddup2(posix_spawn_file_actions_t *
28995 file_actions, int fildes, int newfildes);
```

28996 **DESCRIPTION**

28997 A spawn file actions object is as defined in *posix\_spawn\_file\_actions\_addclose()*.  
 28998 The *posix\_spawn\_file\_actions\_adddup2()* function adds a *dup2()* action to the object referenced by  
 28999 *file\_actions* that will cause the file descriptor *fildes* to be duplicated as *newfildes* (as if *dup2(fildes,*  
 29000 *newfildes)* had been called) when a new process is spawned using this file actions object.

29001 **RETURN VALUE**

29002 Upon successful completion, the *posix\_spawn\_file\_actions\_adddup2()* function shall return zero;  
 29003 otherwise, an error number shall be returned to indicate the error.

29004 **ERRORS**

29005 The *posix\_spawn\_file\_actions\_adddup2()* function shall fail if:

- 29006 [EBADF] The value specified by *fildes* or *newfildes* is negative or greater than or equal to  
 29007 {OPEN\_MAX}.
- 29008 [ENOMEM] Insufficient memory exists to add to the spawn file actions object.

29009 The *posix\_spawn\_file\_actions\_adddup2()* function may fail if:

- 29010 [EINVAL] The value specified by *file\_actions* is invalid.

29011 It shall not be considered an error for the *fildes* argument passed to the  
 29012 *posix\_spawn\_file\_actions\_adddup2()* function to specify a file descriptor for which the specified  
 29013 operation could not be performed at the time of the call. Any such error will be detected when  
 29014 the associated file actions object is later used during a *posix\_spawn()* or *posix\_spawnnp()*  
 29015 operation.

29016 **EXAMPLES**

29017 None.

29018 **APPLICATION USAGE**

29019 The *posix\_spawn\_file\_actions\_adddup2()* function is part of the Spawn option and need not be  
 29020 provided on all implementations.

29021 **RATIONALE**

29022 Refer to the RATIONALE in *posix\_spawn\_file\_actions\_addclose()*.

29023 **FUTURE DIRECTIONS**

29024 None.

29025 **SEE ALSO**

29026 *dup()*, *posix\_spawn()*, *posix\_spawn\_file\_actions\_addclose()*, *posix\_spawn\_file\_actions\_destroy()*,  
 29027 *posix\_spawnnp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <spawn.h>

29028 **CHANGE HISTORY**

29029 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29030 IEEE PASC Interpretation 1003.1 #104 is included, noting that the [EBADF] error can apply to the

29031 *newfildes* argument in addition to *fildes*.

29032 **NAME**

29033 posix\_spawn\_file\_actions\_addopen — add open action to spawn file actions object  
29034 (**REALTIME**)

29035 **SYNOPSIS**

```
29036 SPN #include <spawn.h>
29037 int posix_spawn_file_actions_addopen(posix_spawn_file_actions_t *restrict
29038 file_actions, int fildes, const char *restrict path,
29039 int oflag, mode_t mode);
29040
```

29041 **DESCRIPTION**

29042 Refer to *posix\_spawn\_file\_actions\_addclose()*.

29043 **NAME**

29044 posix\_spawn\_file\_actions\_destroy, posix\_spawn\_file\_actions\_init — destroy and initialize  
 29045 spawn file actions object (**REALTIME**)

29046 **SYNOPSIS**

29047 SPN `#include <spawn.h>`

29048 `int posix_spawn_file_actions_destroy(posix_spawn_file_actions_t *`  
 29049 `file_actions);`

29050 `int posix_spawn_file_actions_init(posix_spawn_file_actions_t *`  
 29051 `file_actions);`

29052

29053 **DESCRIPTION**

29054 A spawn file actions object is as defined in *posix\_spawn\_file\_actions\_addclose()*.

29055 The *posix\_spawn\_file\_actions\_destroy()* function destroys the object referenced by *file\_actions*; the  
 29056 object becomes, in effect, uninitialized. An implementation may cause  
 29057 *posix\_spawn\_file\_actions\_destroy()* to set the object referenced by *file\_actions* to an invalid value. A  
 29058 destroyed spawn file actions object can be reinitialized using *posix\_spawn\_file\_actions\_init()*; the  
 29059 results of otherwise referencing the object after it has been destroyed are undefined.

29060 The *posix\_spawn\_file\_actions\_init()* function initializes the object referenced by *file\_actions* to  
 29061 contain no file actions for *posix\_spawn()* or *posix\_spawnnp()* to perform.

29062 The effect of initializing an already initialized spawn file actions object is undefined.

29063 **RETURN VALUE**

29064 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
 29065 be returned to indicate the error.

29066 **ERRORS**

29067 The *posix\_spawn\_file\_actions\_init()* function shall fail if:

29068 [ENOMEM] Insufficient memory exists to initialize the spawn file actions object.

29069 The *posix\_spawn\_file\_actions\_destroy()* function may fail if:

29070 [EINVAL] The value specified by *file\_actions* is invalid.

29071 **EXAMPLES**

29072 None.

29073 **APPLICATION USAGE**

29074 These functions are part of the Spawn option and need not be provided on all implementations.

29075 **RATIONALE**

29076 Refer to the RATIONALE in *posix\_spawn\_file\_actions\_addclose()*.

29077 **FUTURE DIRECTIONS**

29078 None.

29079 **SEE ALSO**

29080 *posix\_spawn()*, *posix\_spawnnp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<spawn.h>`

29081 **CHANGE HISTORY**

29082 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29083 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

29084 **NAME**

29085        `posix_spawn_file_actions_init` — initialize spawn file actions object (**REALTIME**)

29086 **SYNOPSIS**

29087 SPN        `#include <spawn.h>`

```
29088 int posix_spawn_file_actions_init(posix_spawn_file_actions_t *
29089 file_actions);
```

29090

29091 **DESCRIPTION**

29092        Refer to `posix_spawn_file_actions_destroy()`.

29093 **NAME**

29094 `posix_spawnattr_destroy`, `posix_spawnattr_init` — destroy and initialize spawn attributes object  
 29095 (**REALTIME**)

29096 **SYNOPSIS**

29097 SPN `#include <spawn.h>`

29098 `int posix_spawnattr_destroy(posix_spawnattr_t *attr);`

29099 `int posix_spawnattr_init(posix_spawnattr_t *attr);`

29100

29101 **DESCRIPTION**

29102 A spawn attributes object is of type **posix\_spawnattr\_t** (defined in `<spawn.h>`) and is used to  
 29103 specify the inheritance of process attributes across a spawn operation. IEEE Std. 1003.1-200x  
 29104 does not define comparison or assignment operators for the type **posix\_spawnattr\_t**.

29105 The `posix_spawnattr_destroy()` function destroys a spawn attributes object. The effect of  
 29106 subsequent use of the object is undefined until the object is reinitialized by another call to  
 29107 `posix_spawnattr_init()`. An implementation may cause `posix_spawnattr_destroy()` to set the object  
 29108 referenced by `attr` to an invalid value.

29109 The `posix_spawnattr_init()` function initializes a spawn attributes object `attr` with the default  
 29110 value for all of the individual attributes used by the implementation. The effect of initializing an  
 29111 already initialized spawn attributes option is undefined.

29112 Each implementation shall document the individual attributes it uses and their default values  
 29113 unless these values are defined by IEEE Std. 1003.1-200x. Attributes not defined by  
 29114 IEEE Std. 1003.1-200x, their default values, and the names of the associated functions to get and  
 29115 set those attribute values are implementation-defined.

29116 The resulting spawn attributes object (possibly modified by setting individual attribute values),  
 29117 is used to modify the behavior of `posix_spawn()` or `posix_spawnp()`. After a spawn attributes  
 29118 object has been used to spawn a process by a call to a `posix_spawn()` or `posix_spawnp()`, any  
 29119 function affecting the attributes object (including destruction) does not affect any process that  
 29120 has been spawned in this way.

29121 **RETURN VALUE**

29122 Upon successful completion, `posix_spawnattr_destroy()` and `posix_spawnattr_init()` shall return  
 29123 zero; otherwise, an error number shall be returned to indicate the error.

29124 **ERRORS**

29125 The `posix_spawnattr_init()` function shall fail if:

29126 [ENOMEM] Insufficient memory exists to initialize the spawn attributes object.

29127 The `posix_spawnattr_destroy()` function may fail if:

29128 [EINVAL] The value specified by `attr` is invalid.

29129 **EXAMPLES**

29130 None.

29131 **APPLICATION USAGE**

29132 These functions are part of the Spawn option and need not be provided on all implementations.

29133 **RATIONALE**

29134 The original spawn interface proposed in IEEE Std. 1003.1-200x defined the attributes that  
 29135 specify the inheritance of process attributes across a spawn operation as a structure. In order to  
 29136 be able to separate optional individual attributes under their appropriate options (that is, the  
 29137 `spawn-schedparam` and `spawn-schedpolicy` attributes depending upon the Process Scheduling

29138 option), and also for extensibility and consistency with the newer POSIX interfaces, the  
 29139 attributes interface has been changed to an opaque data type. This interface now consists of the  
 29140 type **posix\_spawnattr\_t**, representing a spawn attributes object, together with associated  
 29141 functions to initialize or destroy the attributes object, and to set or get each individual attribute.  
 29142 Although the new object-oriented interface is more verbose than the original structure, it is  
 29143 simple to use, more extensible, and easy to implement.

29144 **FUTURE DIRECTIONS**

29145 None.

29146 **SEE ALSO**

29147 *posix\_spawn()*, *posix\_spawnattr\_getsigdefault()*, *posix\_spawnattr\_getflags()*,  
 29148 *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*,  
 29149 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,  
 29150 *posix\_spawnattr\_setpgroup()*, *posix\_spawnattr\_setsigmask()*, *posix\_spawnattr\_setschedpolicy()*,  
 29151 *posix\_spawnattr\_setschedparam()*, *posix\_spawnnp()*, the Base Definitions volume of  
 29152 IEEE Std. 1003.1-200x, <spawn.h>

29153 **CHANGE HISTORY**

29154 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29155 IEEE PASC Interpretation 1003.1 #106 is included, noting that the effect of initializing an already  
 29156 initialized spawn attributes option is undefined.

29157 **NAME**

29158 posix\_spawnattr\_getflags, posix\_spawnattr\_setflags — get and set spawn-flags attribute of  
 29159 spawn attributes object (**REALTIME**)

29160 **SYNOPSIS**

29161 SPN `#include <spawn.h>`

29162 `int posix_spawnattr_getflags(const posix_spawnattr_t *restrict attr,`  
 29163 `short *restrict flags);`

29164 `int posix_spawnattr_setflags(posix_spawnattr_t *attr, short flags);`  
 29165

29166 **DESCRIPTION**

29167 The *spawn-flags* attribute is used to indicate which process attributes are to be changed in the  
 29168 new process image when invoking *posix\_spawn()* or *posix\_spawnp()*. It is the bitwise-inclusive  
 29169 OR of zero or more of the flags POSIX\_SPAWN\_RESETIDS, POSIX\_SPAWN\_SETPGROUP,  
 29170 PS POSIX\_SPAWN\_SETSIGDEF, and POSIX\_SPAWN\_SETSIGMASK,  
 29171 POSIX\_SPAWN\_SETSCHEDPARAM, and POSIX\_SPAWN\_SETSCHEDULER. In addition, if the  
 29172 Process Scheduling option is supported, the flags POSIX\_SPAWN\_SETSCHEDPARAM and  
 29173 POSIX\_SPAWN\_SETSCHEDULER shall also be supported. These flags are defined in  
 29174 `<spawn.h>`. The default value of this attribute shall be with no flags set.

29175 The *posix\_spawnattr\_getflags()* function obtains the value of the *spawn-flags* attribute from the  
 29176 attributes object referenced by *attr*.

29177 The *posix\_spawnattr\_setflags()* function is used to set the *spawn-flags* attribute in an initialized  
 29178 attributes object referenced by *attr*.

29179 **RETURN VALUE**

29180 Upon successful completion, *posix\_spawnattr\_getflags()* shall return zero and store the value of  
 29181 the *spawn-flags* attribute of *attr* into the object referenced by the *flags* parameter; otherwise, an  
 29182 error number shall be returned to indicate the error.

29183 Upon successful completion, *posix\_spawnattr\_setflags()* shall return zero; otherwise, an error  
 29184 number shall be returned to indicate the error.

29185 **ERRORS**

29186 These functions may fail if:

29187 [EINVAL] The value specified by *attr* is invalid.

29188 The *posix\_spawnattr\_setflags()* function may fail if:

29189 [EINVAL] The value of the attribute being set is not valid.

29190 **EXAMPLES**

29191 None.

29192 **APPLICATION USAGE**

29193 These functions are part of the Spawn option and need not be provided on all implementations.

29194 **RATIONALE**

29195 None.

29196 **FUTURE DIRECTIONS**

29197 None.

29198 **SEE ALSO**

29199 *posix\_spawn()*, *posix\_spawnattr\_destroy()* (on page 1405), *posix\_spawnattr\_init()* (on page 1419), *posix\_spawnattr\_getsigdefault()*,  
29200 *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*,  
29201 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setpgroup()*,  
29202 *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
29203 *posix\_spawnnp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <spawn.h>

29205 **CHANGE HISTORY**

29206 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29207 **NAME**

29208 posix\_spawnattr\_getpgroup, posix\_spawnattr\_setpgroup — get and set spawn-pgroup attribute  
 29209 of spawn attributes object (**REALTIME**)

29210 **SYNOPSIS**

```
29211 SPN #include <spawn.h>
29212
29212 int posix_spawnattr_getpgroup(const posix_spawnattr_t *restrict attr,
29213 pid_t *restrict pgroup);
29214 int posix_spawnattr_setpgroup(posix_spawnattr_t *attr, pid_t pgroup);
29215
```

29216 **DESCRIPTION**

29217 The *spawn-pgroup* attribute represents the process group to be joined by the new process image  
 29218 in a spawn operation (if `POSIX_SPAWN_SETPGROUP` is set in the *spawn-flags* attribute). The  
 29219 default value of this attribute shall be zero.

29220 The *posix\_spawnattr\_getpgroup()* function obtains the value of the *spawn-pgroup* attribute from  
 29221 the attributes object referenced by *attr*.

29222 The *posix\_spawnattr\_setpgroup()* function is used to set the *spawn-pgroup* attribute in an  
 29223 initialized attributes object referenced by *attr*.

29224 **RETURN VALUE**

29225 Upon successful completion, *posix\_spawnattr\_getpgroup()* shall return zero and store the value of  
 29226 the *spawn-pgroup* attribute of *attr* into the object referenced by the *pgroup* parameter; otherwise,  
 29227 an error number shall be returned to indicate the error.

29228 Upon successful completion, *posix\_spawnattr\_setpgroup()* shall return zero; otherwise, an error  
 29229 number shall be returned to indicate the error.

29230 **ERRORS**

29231 These functions may fail if:

29232 [EINVAL] The value specified by *attr* is invalid.

29233 The *posix\_spawnattr\_setpgroup()* function may fail if:

29234 [EINVAL] The value of the attribute being set is not valid.

29235 **EXAMPLES**

29236 None.

29237 **APPLICATION USAGE**

29238 These functions are part of the Spawn option and need not be provided on all implementations.

29239 **RATIONALE**

29240 None.

29241 **FUTURE DIRECTIONS**

29242 None.

29243 **SEE ALSO**

29244 *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getsigdefault()*,  
 29245 *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*,  
 29246 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,  
 29247 *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
 29248 *posix\_spawnp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <spawn.h>

29249 **CHANGE HISTORY**

29250 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29251 **NAME**

29252 `posix_spawnattr_getschedparam`, `posix_spawnattr_setschedparam` — get and set spawn-  
 29253 `schedparam` attribute of spawn attributes object (**REALTIME**)

29254 **SYNOPSIS**

```
29255 SPN PS #include <spawn.h>
29256 #include <sched.h>

29257 int posix_spawnattr_getschedparam(const posix_spawnattr_t *restrict attr,
29258 struct sched_param *restrict schedparam);
29259 int posix_spawnattr_setschedparam(posix_spawnattr_t *restrict attr,
29260 const struct sched_param *restrict schedparam);
29261
```

29262 **DESCRIPTION**

29263 The *spawn-schedparam* attribute represents the scheduling parameters to be assigned to the new  
 29264 process image in a spawn operation (if `POSIX_SPAWN_SETSCHEDULER` or  
 29265 `POSIX_SPAWN_SETSCHEDPARAM` is set in the *spawn-flags* attribute). The default value of this  
 29266 attribute is unspecified.

29267 The *posix\_spawnattr\_getschedparam()* function obtains the value of the *spawn-schedparam* attribute  
 29268 from the attributes object referenced by *attr*.

29269 The *posix\_spawnattr\_setschedparam()* function is used to set the *spawn-schedparam* attribute in an  
 29270 initialized attributes object referenced by *attr*.

29271 **RETURN VALUE**

29272 Upon successful completion, *posix\_spawnattr\_getschedparam()* shall return zero and store the  
 29273 value of the *spawn-schedparam* attribute of *attr* into the object referenced by the *schedparam*  
 29274 parameter; otherwise, an error number shall be returned to indicate the error.

29275 Upon successful completion, *posix\_spawnattr\_setschedparam()* shall return zero; otherwise, an  
 29276 error number shall be returned to indicate the error.

29277 **ERRORS**

29278 These functions may fail if:

29279 [EINVAL] The value specified by *attr* is invalid.

29280 The *posix\_spawnattr\_setschedparam()* function may fail if:

29281 [EINVAL] The value of the attribute being set is not valid.

29282 **EXAMPLES**

29283 None.

29284 **APPLICATION USAGE**

29285 These functions are part of the Spawn and Process Scheduling options and need not be provided  
 29286 on all implementations.

29287 **RATIONALE**

29288 None.

29289 **FUTURE DIRECTIONS**

29290 None.

29291 **SEE ALSO**

29292 *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getsigdefault()*,  
 29293 *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedpolicy()*,  
 29294 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,

29295 *posix\_spawnattr\_setpgroup()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
29296 *posix\_spawnnp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sched.h>, <spawn.h>

29297 **CHANGE HISTORY**

29298 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29299 **NAME**

29300 `posix_spawnattr_getschedpolicy`, `posix_spawnattr_setschedpolicy` — get and set spawn-  
 29301 `schedpolicy` attribute of spawn attributes object (**REALTIME**)

29302 **SYNOPSIS**

```
29303 SPN PS #include <spawn.h>
29304 #include <sched.h>

29305 int posix_spawnattr_getschedpolicy(const posix_spawnattr_t *restrict attr,
29306 int *restrict schedpolicy);
29307 int posix_spawnattr_setschedpolicy(posix_spawnattr_t *attr,
29308 int schedpolicy);
29309
```

29310 **DESCRIPTION**

29311 The *spawn-schedpolicy* attribute represents the scheduling policy to be assigned to the new  
 29312 process image in a spawn operation (if `POSIX_SPAWN_SETSCHEDULER` is set in the *spawn-*  
 29313 *flags* attribute). The default value of this attribute is unspecified.

29314 The *posix\_spawnattr\_getschedpolicy()* function obtains the value of the *spawn-schedpolicy* attribute  
 29315 from the attributes object referenced by *attr*.

29316 The *posix\_spawnattr\_setschedpolicy()* function is used to set the *spawn-schedpolicy* attribute in an  
 29317 initialized attributes object referenced by *attr*.

29318 **RETURN VALUE**

29319 Upon successful completion, *posix\_spawnattr\_getschedpolicy()* shall return zero and store the  
 29320 value of the *spawn-schedpolicy* attribute of *attr* into the object referenced by the *schedpolicy*  
 29321 parameter; otherwise, an error number shall be returned to indicate the error.

29322 Upon successful completion, *posix\_spawnattr\_setschedpolicy()* shall return zero; otherwise, an  
 29323 error number shall be returned to indicate the error.

29324 **ERRORS**

29325 These functions may fail if:

29326 [EINVAL] The value specified by *attr* is invalid.

29327 The *posix\_spawnattr\_setschedpolicy()* function may fail if:

29328 [EINVAL] The value of the attribute being set is not valid.

29329 **EXAMPLES**

29330 None.

29331 **APPLICATION USAGE**

29332 These functions are part of the Spawn and Process Scheduling options and need not be provided  
 29333 on all implementations.

29334 **RATIONALE**

29335 None.

29336 **FUTURE DIRECTIONS**

29337 None.

29338 **SEE ALSO**

29339 *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getsigdefault()*,  
 29340 *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedparam()*,  
 29341 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,  
 29342 *posix\_spawnattr\_setpgroup()*, *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setsigmask()*,

- 29343            *posix\_spawnp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sched.h>, <spawn.h>
- 29344 **CHANGE HISTORY**
- 29345            First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29346 **NAME**

29347 posix\_spawnattr\_getsigdefault, posix\_spawnattr\_setsigdefault — get and set spawn-sigdefault  
 29348 attribute of spawn attributes object (**REALTIME**)

29349 **SYNOPSIS**

```
29350 SPN #include <signal.h>
29351 #include <spawn.h>

29352 int posix_spawnattr_getsigdefault(const posix_spawnattr_t *restrict attr,
29353 sigset_t *restrict sigdefault);
29354 int posix_spawnattr_setsigdefault(posix_spawnattr_t *restrict attr,
29355 const sigset_t *restrict sigdefault);
29356
```

29357 **DESCRIPTION**

29358 The *spawn-sigdefault* attribute represents the set of signals to be forced to default signal handling  
 29359 in the new process image (if POSIX\_SPAWN\_SETSIGDEF is set in the *spawn-flags* attribute) by a  
 29360 spawn operation. The default value of this attribute shall be an empty signal set.

29361 The *posix\_spawnattr\_getsigdefault()* function obtains the value of the *spawn-sigdefault* attribute  
 29362 from the attributes object referenced by *attr*.

29363 The *posix\_spawnattr\_setsigdefault()* function is used to set the *spawn-sigdefault* attribute in an  
 29364 initialized attributes object referenced by *attr*.

29365 **RETURN VALUE**

29366 Upon successful completion, *posix\_spawnattr\_getsigdefault()* shall return zero and store the value  
 29367 of the *spawn-sigdefault* attribute of *attr* into the object referenced by the *sigdefault* parameter;  
 29368 otherwise, an error number shall be returned to indicate the error.

29369 Upon successful completion, *posix\_spawnattr\_setsigdefault()* shall return zero; otherwise, an error  
 29370 number shall be returned to indicate the error.

29371 **ERRORS**

29372 These functions may fail if:

29373 [EINVAL] The value specified by *attr* is invalid.

29374 The *posix\_spawnattr\_setsigdefault()* function may fail if:

29375 [EINVAL] The value of the attribute being set is not valid.

29376 **EXAMPLES**

29377 None.

29378 **APPLICATION USAGE**

29379 These functions are part of the Spawn option and need not be provided on all implementations.

29380 **RATIONALE**

29381 None.

29382 **FUTURE DIRECTIONS**

29383 None.

29384 **SEE ALSO**

29385 *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getflags()*,  
 29386 *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedparam()*, *posix\_spawnattr\_getschedpolicy()*,  
 29387 *posix\_spawnattr\_getsigmask()*, *posix\_spawnattr\_setflags()*, *posix\_spawnattr\_setpgroup()*,  
 29388 *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*, *posix\_spawnattr\_setsigmask()*,  
 29389 *posix\_spawnnp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <signal.h>, <spawn.h>

29390 **CHANGE HISTORY**

29391 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29392 **NAME**

29393 posix\_spawnattr\_getsigmask, posix\_spawnattr\_setsigmask — get and set spawn-sigmask  
 29394 attribute of spawn attributes object (**REALTIME**)

29395 **SYNOPSIS**

29396 SPN #include <signal.h>

29397 #include <spawn.h>

```
29398 int posix_spawnattr_getsigmask(const posix_spawnattr_t *restrict attr,
29399 sigset_t *restrict sigmask);
```

```
29400 int posix_spawnattr_setsigmask(posix_spawnattr_t *restrict attr,
29401 const sigset_t *restrict sigmask);
```

29402

29403 **DESCRIPTION**

29404 The *spawn-sigmask* attribute represents the signal mask in effect in the new process image of a  
 29405 spawn operation (if `POSIX_SPAWN_SETSIGMASK` is set in the *spawn-flags* attribute). The  
 29406 default value of this attribute is unspecified.

29407 The *posix\_spawnattr\_getsigmask()* function obtains the value of the *spawn-sigmask* attribute from  
 29408 the attributes object referenced by *attr*.

29409 The *posix\_spawnattr\_setsigmask()* function is used to set the *spawn-sigmask* attribute in an  
 29410 initialized attributes object referenced by *attr*.

29411 **RETURN VALUE**

29412 Upon successful completion, *posix\_spawnattr\_getsigmask()* shall return zero and store the value  
 29413 of the *spawn-sigmask* attribute of *attr* into the object referenced by the *sigmask* parameter;  
 29414 otherwise, an error number shall be returned to indicate the error.

29415 Upon successful completion, *posix\_spawnattr\_setsigmask()* shall return zero; otherwise, an error  
 29416 number shall be returned to indicate the error.

29417 **ERRORS**

29418 These functions may fail if:

29419 [EINVAL] The value specified by *attr* is invalid.

29420 The *posix\_spawnattr\_setsigmask()* function may fail if:

29421 [EINVAL] The value of the attribute being set is not valid.

29422 **EXAMPLES**

29423 None.

29424 **APPLICATION USAGE**

29425 These functions are part of the Spawn option and need not be provided on all implementations.

29426 **RATIONALE**

29427 None.

29428 **FUTURE DIRECTIONS**

29429 None.

29430 **SEE ALSO**

29431 *posix\_spawn()*, *posix\_spawnattr\_destroy()*, *posix\_spawnattr\_init()*, *posix\_spawnattr\_getsigdefault()*,  
 29432 *posix\_spawnattr\_getflags()*, *posix\_spawnattr\_getpgroup()*, *posix\_spawnattr\_getschedparam()*,  
 29433 *posix\_spawnattr\_getschedpolicy()*, *posix\_spawnattr\_setsigdefault()*, *posix\_spawnattr\_setflags()*,  
 29434 *posix\_spawnattr\_setpgroup()*, *posix\_spawnattr\_setschedparam()*, *posix\_spawnattr\_setschedpolicy()*,  
 29435 *posix\_spawnnp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <signal.h>, <spawn.h>

29436 **CHANGE HISTORY**

29437 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

29438 **NAME**29439 `posix_spawnattr_init` — initialize spawn attributes object (**REALTIME**)29440 **SYNOPSIS**29441 SPN `#include <spawn.h>`29442 `int posix_spawnattr_init(posix_spawnattr_t *attr);`

29443

29444 **DESCRIPTION**29445 Refer to `posix_spawnattr_destroy()`.

29446 **NAME**

29447        `posix_spawnattr_setflags` — set spawn-flags attribute of spawn attributes object (**REALTIME**)

29448 **SYNOPSIS**

29449 SPN    `#include <spawn.h>`

29450        `int posix_spawnattr_setflags(posix_spawnattr_t *attr, short flags);`

29451

29452 **DESCRIPTION**

29453        Refer to `posix_spawnattr_getflags()`.

29454 **NAME**

29455        `posix_spawnattr_setpgroup` — set spawn-pgroup attribute of spawn attributes object  
29456        (**REALTIME**)

29457 **SYNOPSIS**

29458 SPN    `#include <spawn.h>`

29459        `int posix_spawnattr_setpgroup(posix_spawnattr_t *attr, pid_t pgroup);`

29460

29461 **DESCRIPTION**

29462        Refer to `posix_spawnattr_getpgroup()`.

29463 **NAME**

29464 posix\_spawnattr\_setschedparam — set spawn-schedparam attribute of spawn attributes object  
29465 (**REALTIME**)

29466 **SYNOPSIS**

29467 SPN PS #include <sched.h>

29468 #include <spawn.h>

29469 int posix\_spawnattr\_setschedparam(posix\_spawnattr\_t \*restrict attr,  
29470 const struct sched\_param \*restrict schedparam);

29471

29472 **DESCRIPTION**

29473 Refer to *posix\_spawnattr\_getschedparam()*.

29474 **NAME**

29475 `posix_spawnattr_setschedpolicy` — set spawn-schedpolicy attribute of spawn attributes object  
29476 **(REALTIME)**

29477 **SYNOPSIS**

29478 SPN PS `#include <sched.h>`

29479 `#include <spawn.h>`

```
29480 int posix_spawnattr_setschedpolicy(posix_spawnattr_t *attr,
29481 int schedpolicy);
29482
```

29483 **DESCRIPTION**

29484 Refer to `posix_spawnattr_getschedpolicy()`.

29485 **NAME**

29486        posix\_spawnattr\_setsigdefault — set spawn-sigdefault attribute of spawn attributes object  
29487        (**REALTIME**)

29488 **SYNOPSIS**

29489 SPN    #include <signal.h>

29490        #include <spawn.h>

29491        int posix\_spawnattr\_setsigdefault(posix\_spawnattr\_t \*restrict attr,  
29492                                        const sigset\_t \*restrict sigdefault);

29493

29494 **DESCRIPTION**

29495        Refer to *posix\_spawnattr\_getsigdefault()*.

29496 **NAME**

29497 posix\_spawnattr\_setsigmask — set spawn-sigmask attribute of spawn attributes object  
29498 (**REALTIME**)

29499 **SYNOPSIS**

29500 SPN #include <signal.h>

29501 #include <spawn.h>

29502 int posix\_spawnattr\_setsigmask(posix\_spawnattr\_t \*restrict attr,  
29503 const sigset\_t \*restrict sigmask);

29504

29505 **DESCRIPTION**

29506 Refer to *posix\_spawnattr\_getsigmask()*.

29507 **NAME**

29508        posix\_spawnnp — spawn a process (**REALTIME**)

29509 **SYNOPSIS**

29510 SPN     #include <spawn.h>

```
29511 int posix_spawnnp(pid_t *restrict pid, const char *restrict file,
29512 const posix_spawn_file_actions_t *file_actions,
29513 const posix_spawnattr_t *restrict attrp,
29514 char *const argv[restrict], char *const envp[restrict]);
29515
```

29516 **DESCRIPTION**

29517        Refer to *posix\_spawn()*.

29518 **NAME**

29519        posix\_trace\_attr\_destroy, posix\_trace\_attr\_init — trace stream attributes object destroy and  
29520        initialization

29521 **SYNOPSIS**

29522 TRC     #include <trace.h>

29523        int posix\_trace\_attr\_destroy(trace\_attr\_t \*attr);

29524        int posix\_trace\_attr\_init(trace\_attr\_t \*attr);

29525

29526 **DESCRIPTION**

29527        The *posix\_trace\_attr\_destroy()* function is used to destroy an initialized trace attributes object.  
29528        The results of using the attributes object after it has been destroyed are unspecified. A destroyed  
29529        trace attributes object can be reinitialized using *posix\_trace\_attr\_init()*.

29530        The *posix\_trace\_attr\_init()* function initializes a trace attributes object *attr* with the default value  
29531        for all of the individual attributes used by a given implementation. The read-only *generation-*  
29532        *version* and *clock-resolution* attributes of the newly initialized trace attributes object shall be set to  
29533        their appropriate values (Section 2.11.1.2 (on page 578)).

29534        The effect of initializing an already-initialized trace attributes object is unspecified.

29535        Implementations may add extensions to the trace attributes object structure as permitted in the  
29536        Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 2, Conformance.

29537        The resulting attributes object (possibly modified by setting individual attributes values), when  
29538        used by *posix\_trace\_create()*, defines the attributes of the trace stream created. A single attributes  
29539        object can be used in multiple calls to *posix\_trace\_create()*. After one or more trace streams have  
29540        been created using an attributes object, any function affecting that attributes object, including  
29541        destruction, does not affect any trace stream previously created. An initialized attributes object  
29542        also serves to receive the attributes of an existing trace stream or trace log when calling the  
29543        *posix\_trace\_get\_attr()* function.

29544 **RETURN VALUE**

29545        Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
29546        return the corresponding error number.

29547 **ERRORS**

29548        The *posix\_trace\_attr\_destroy()* function may fail if:

29549        [EINVAL]        The value of *attr* is invalid.

29550        The *posix\_trace\_attr\_init()* function shall fail if:

29551        [ENOMEM]       Insufficient memory exists to initialize the trace attributes object.

29552 **EXAMPLES**

29553        None.

29554 **APPLICATION USAGE**

29555        None.

29556 **RATIONALE**

29557        None.

29558 **FUTURE DIRECTIONS**

29559        None.

29560 **SEE ALSO**

29561            *posix\_trace\_create()*, *posix\_trace\_get\_attr()*, *uname()*, the Base Definitions volume of  
29562            IEEE Std. 1003.1-200x, <trace.h>

29563 **CHANGE HISTORY**

29564            First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

29565 **NAME**

29566 `posix_trace_attr_getclockres`, `posix_trace_attr_getcreatetime`, `posix_trace_attr_getgenversion`,  
 29567 `posix_trace_attr_getname`, `posix_trace_attr_setname` — retrieve and set information about a  
 29568 trace stream

29569 **SYNOPSIS**

```
29570 TRC #include <time.h>
29571 #include <trace.h>

29572 int posix_trace_attr_getclockres(const trace_attr_t *attr,
29573 struct timespec *resolution);
29574 int posix_trace_attr_getcreatetime(const trace_attr_t *attr,
29575 struct timespec *createtime);

29576 #include <trace.h>

29577 int posix_trace_attr_getgenversion(const trace_attr_t *attr,
29578 char *genversion);
29579 int posix_trace_attr_getname(const trace_attr_t *attr,
29580 char *tracename);
29581 int posix_trace_attr_setname(trace_attr_t *attr,
29582 const char *tracename);
29583
```

29584 **DESCRIPTION**

29585 The `posix_trace_attr_getclockres()` function shall copy the clock resolution of the clock used to  
 29586 generate timestamps from the *clock-resolution* attribute of the attributes object pointed to by the  
 29587 *attr* argument into the structure pointed to by the *resolution* argument.

29588 The `posix_trace_attr_getcreatetime()` function shall copy the trace stream creation time from the  
 29589 *creation-time* attribute of the attributes object pointed to by the *attr* argument into the structure  
 29590 pointed to by the *createtime* argument. The *creation-time* attribute shall represent the time of  
 29591 creation of the trace stream.

29592 The `posix_trace_attr_getgenversion()` function shall copy the string containing version information  
 29593 from the *generation-version* attribute of the attributes object pointed to by the *attr* argument into  
 29594 the string pointed to by the *genversion* argument. The *genversion* argument shall be the address of  
 29595 a character array which can store at least {TRACE\_NAME\_MAX} characters.

29596 The `posix_trace_attr_getname()` function shall copy the string containing the trace name from the  
 29597 *trace-name* attribute of the attributes object pointed to by the *attr* argument into the string  
 29598 pointed to by the *tracename* argument. The *tracename* argument shall be the address of a character  
 29599 array which can store at least {TRACE\_NAME\_MAX} characters.

29600 The `posix_trace_attr_setname()` function shall set the name in the *trace-name* attribute of the  
 29601 attributes object pointed to by the *attr* argument, using the trace name string supplied by the  
 29602 *tracename* argument. If the supplied string contains more than {TRACE\_NAME\_MAX}  
 29603 characters, the name copied into the *trace-name* attribute may be truncated to one less than the  
 29604 length of {TRACE\_NAME\_MAX} characters. The default value is a null string.

29605 **RETURN VALUE**

29606 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 29607 return the corresponding error number.

29608 If successful, the `posix_trace_attr_getclockres()` function stores the *clock-resolution* attribute value  
 29609 in the object pointed to by *resolution*. Otherwise, the content of this object is unspecified.

29610 If successful, the *posix\_trace\_attr\_getcreatetime()* function stores the trace stream creation time in  
29611 the object pointed to by *createtime*. Otherwise, the content of this object is unspecified.

29612 If successful, the *posix\_trace\_attr\_getgenversion()* function stores the trace version information in  
29613 the string pointed to by *genversion*. Otherwise, the content of this string is unspecified.

29614 If successful, the *posix\_trace\_attr\_getname()* function stores the trace name in the string pointed  
29615 to by *tracename*. Otherwise, the content of this string is unspecified.

## 29616 ERRORS

29617 The *posix\_trace\_attr\_getclockres()*, *posix\_trace\_attr\_getcreatetime()*, *posix\_trace\_attr\_getgenversion()*,  
29618 and *posix\_trace\_attr\_getname()* functions may fail if:

29619 [EINVAL] The value specified by one of the arguments is invalid.

## 29620 EXAMPLES

29621 None.

## 29622 APPLICATION USAGE

29623 None.

## 29624 RATIONALE

29625 None.

## 29626 FUTURE DIRECTIONS

29627 None.

## 29628 SEE ALSO

29629 *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_get\_attr()*, *uname()*, the Base Definitions  
29630 volume of IEEE Std. 1003.1-200x, <time.h>, <trace.h>

## 29631 CHANGE HISTORY

29632 First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

## 29633 NAME

29634 posix\_trace\_attr\_getinherited, posix\_trace\_attr\_getlogfullpolicy,  
 29635 posix\_trace\_attr\_getstreamfullpolicy, posix\_trace\_attr\_setinherited,  
 29636 posix\_trace\_attr\_setlogfullpolicy, posix\_trace\_attr\_setstreamfullpolicy — retrieve and set the  
 29637 behavior of a trace stream

## 29638 SYNOPSIS

```
29639 TRC #include <trace.h>
29640 TRC TRI int posix_trace_attr_getinherited(const trace_attr_t *attr,
29641 int *inheritancepolicy);
29642 TRC TRL int posix_trace_attr_getlogfullpolicy(const trace_attr_t *attr,
29643 int *logpolicy);
29644 TRC int posix_trace_attr_getstreamfullpolicy(const trace_attr_t *attr,
29645 int *streampolicy);
29646 TRC TRI int posix_trace_attr_setinherited(trace_attr_t *attr,
29647 int inheritancepolicy);
29648 TRC TRL int posix_trace_attr_setlogfullpolicy(trace_attr_t *attr,
29649 int logpolicy);
29650 TRC int posix_trace_attr_setstreamfullpolicy(trace_attr_t *attr,
29651 int streampolicy);
29652
```

## 29653 DESCRIPTION

29654 TRI The *posix\_trace\_attr\_getinherited()* and *posix\_trace\_attr\_setinherited()* functions, respectively, get  
 29655 and set the inheritance policy stored in the *inheritance* attribute for traced processes across the  
 29656 *fork()* and *spawn()* operations. The *inheritance* attribute of the attributes object pointed to by the  
 29657 *attr* argument shall be set to one of the following values defined by manifest constants in the  
 29658 **<trace.h>** header:

## 29659 POSIX\_TRACE\_CLOSE\_FOR\_CHILD

29660 After a *fork()* or *spawn()* operation, the child shall not be traced, and tracing of the parent  
 29661 shall continue.

## 29662 POSIX\_TRACE\_INHERITED

29663 After a *fork()* or *spawn()* operation, if the parent is being traced, its child shall be  
 29664 concurrently traced using the same trace stream.

29665 The default value for the *inheritance* attribute is **POSIX\_TRACE\_CLOSE\_FOR\_CHILD**.

29666 TRL The *posix\_trace\_attr\_getlogfullpolicy()* and *posix\_trace\_attr\_setlogfullpolicy()* functions,  
 29667 respectively, get and set the trace log full policy stored in the *log-full-policy* attribute of the  
 29668 attributes object pointed to by the *attr* argument.

29669 The *log-full-policy* attribute shall be set to one of the following values defined by manifest  
 29670 constants in the **<trace.h>** header:

## 29671 POSIX\_TRACE\_LOOP

29672 The trace log shall loop until the associated trace stream is stopped. This policy means that  
 29673 when the trace log gets full, the file system shall reuse the resources allocated to the oldest  
 29674 trace events that were recorded. In this way, the trace log will always contain the most  
 29675 recent trace events flushed.

## 29676 POSIX\_TRACE\_UNTIL\_FULL

29677 The trace stream shall be flushed to the trace log until the trace log is full. This condition can  
 29678 be deduced from the *posix\_log\_full\_status* member status (see the *posix\_trace\_status\_info()*  
 29679 function). The last recorded trace event shall be the **POSIX\_TRACE\_STOP** trace event.

29680 POSIX\_TRACE\_APPEND  
 29681 The associated trace stream shall be flushed to the trace log without log size limitation. If  
 29682 the application specifies POSIX\_TRACE\_APPEND, the implementation shall ignore the  
 29683 *log-max-size* attribute.

29684 The default value for the *log-full-policy* attribute is POSIX\_TRACE\_LOOP.

29685 The *posix\_trace\_attr\_getstreamfullpolicy()* and *posix\_trace\_attr\_setstreamfullpolicy()* functions,  
 29686 respectively, get and set the trace stream full policy stored in the *stream-full-policy* attribute of the  
 29687 attributes object pointed to by the *attr* argument.

29688 The *stream-full-policy* attribute shall be set to one of the following values defined by manifest  
 29689 constants in the <trace.h> header:

29690 POSIX\_TRACE\_LOOP  
 29691 The trace stream shall loop until explicitly stopped by the *posix\_trace\_stop()* function. This  
 29692 policy means that when the trace stream is full, the trace system shall reuse the resources  
 29693 allocated to the oldest trace events recorded. In this way, the trace stream will always  
 29694 contain the most recent trace events recorded.

29695 POSIX\_TRACE\_UNTIL\_FULL  
 29696 The trace stream will run until the trace stream resources are exhausted. Then the trace  
 29697 stream will stop. This condition can be deduced from *posix\_stream\_status* and  
 29698 *posix\_stream\_full\_status* statuses (see the *posix\_trace\_status\_info()* function). When this trace  
 29699 stream is read, a POSIX\_TRACE\_STOP trace event shall be reported after reporting the last  
 29700 recorded trace event. The trace system shall reuse the resources allocated to any trace  
 29701 events already reported—see the *posix\_trace\_getnext\_event()*, *posix\_trace\_trygetnext\_event()*,  
 29702 and *posix\_trace\_timedgetnext\_event()* functions—or already flushed for an active trace stream  
 29703 with log if the Trace Log option is supported; see the *posix\_trace\_flush()* function. The trace  
 29704 system shall restart the trace stream when it is empty and may restart it sooner. A  
 29705 POSIX\_TRACE\_START trace event shall be reported before reporting the next recorded  
 29706 trace event.

29707 POSIX\_TRACE\_FLUSH  
 29708 If the Trace Log option is supported, this policy is identical to the  
 29709 POSIX\_TRACE\_UNTIL\_FULL trace stream full policy except that the trace stream shall be  
 29710 flushed regularly as if *posix\_trace\_flush()* had been explicitly called. Defining this policy for  
 29711 an active trace stream without log shall be invalid.

29712 The default value for the *stream-full-policy* attribute shall be POSIX\_TRACE\_LOOP for an active  
 29713 trace stream without log.

29714 If the Trace Log option is supported, the default value for the *stream-full-policy* attribute shall be  
 29715 POSIX\_TRACE\_FLUSH for an active trace stream with log.

29716 **RETURN VALUE**  
 29717 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 29718 return the corresponding error number.

29719 TRI If successful, the *posix\_trace\_attr\_getinherited()* function stores the *inheritance* attribute value in  
 29720 the object pointed to by *inheritancepolicy*. Otherwise, the content of this object is undefined.

29721 TRL If successful, the *posix\_trace\_attr\_getlogfullpolicy()* function stores the *log-full-policy* attribute  
 29722 value in the object pointed to by *logpolicy*. Otherwise, the content of this object is undefined.

29723 If successful, the *posix\_trace\_attr\_getstreamfullpolicy()* function stores the *stream-full-policy*  
 29724 attribute value in the object pointed to by *streampolicy*. Otherwise, the content of this object is  
 29725 undefined.

29726 **ERRORS**

29727 These functions may fail if:

29728 [EINVAL] The value specified by at least one of the arguments is invalid.

29729 **EXAMPLES**

29730 None.

29731 **APPLICATION USAGE**

29732 None.

29733 **RATIONALE**

29734 None.

29735 **FUTURE DIRECTIONS**

29736 None.

29737 **SEE ALSO**

29738 *fork()*, *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_flush()*, *posix\_trace\_get\_attr()*,  
29739 *posix\_trace\_getnext\_event()*, *posix\_trace\_start()*, **posix\_trace\_status\_info Structure**,  
29740 *posix\_trace\_timedgetnext\_event()*, <REFERENCE UNDEFINED>(spawn), the Base Definitions  
29741 volume of IEEE Std. 1003.1-200x, <trace.h>

29742 **CHANGE HISTORY**

29743 First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

29744 **NAME**

29745 posix\_trace\_attr\_getlogsize, posix\_trace\_attr\_getmaxdatasize,  
 29746 posix\_trace\_attr\_getmaxsystemeventsize, posix\_trace\_attr\_getmaxusereventsize,  
 29747 posix\_trace\_attr\_getstreamsize, posix\_trace\_attr\_setlogsize, posix\_trace\_attr\_setmaxdatasize,  
 29748 posix\_trace\_attr\_setstreamsize — retrieve and set trace stream size attributes

29749 **SYNOPSIS**

```

29750 TRC #include <sys/types.h>
29751 #include <trace.h>

29752 TRC TRL int posix_trace_attr_getlogsize(const trace_attr_t *attr,
29753 size_t *logsize);
29754 TRC int posix_trace_attr_getmaxdatasize(const trace_attr_t *attr,
29755 size_t *maxdatasize);
29756 int posix_trace_attr_getmaxsystemeventsize(const trace_attr_t *attr,
29757 size_t *eventsize);
29758 int posix_trace_attr_getmaxusereventsize(const trace_attr_t *attr,
29759 size_t data_len, size_t *eventsize);
29760 int posix_trace_attr_getstreamsize(const trace_attr_t *attr,
29761 size_t *streamsize);
29762 TRC TRL int posix_trace_attr_setlogsize(trace_attr_t *attr,
29763 size_t logsize);
29764 TRC int posix_trace_attr_setmaxdatasize(trace_attr_t *attr,
29765 size_t maxdatasize);
29766 int posix_trace_attr_setstreamsize(trace_attr_t *attr,
29767 size_t streamsize);
29768

```

29769 **DESCRIPTION**

29770 TRL The *posix\_trace\_attr\_getlogsize()* function shall copy the log size, in bytes, from the *log-max-size*  
 29771 attribute of the attributes object pointed to by the *attr* argument into the variable pointed to by  
 29772 the *logsize* argument. This log size is the maximum total of bytes that shall be allocated for  
 29773 system and user trace events in the trace log. The default value for the *log-max-size* attribute is  
 29774 implementation-defined.

29775 The *posix\_trace\_attr\_setlogsize()* function shall set the maximum allowed size, in bytes, in the  
 29776 *log-max-size* attribute of the attributes object pointed to by the *attr* argument, using the size value  
 29777 supplied by the *logsize* argument.

29778 The trace log size shall be used if the *log-full-policy* attribute is set to *POSIX\_TRACE\_LOOP* or  
 29779 *POSIX\_TRACE\_UNTIL\_FULL*. If the *log-full-policy* attribute is set to *POSIX\_TRACE\_APPEND*,  
 29780 the implementation shall ignore the *log-max-size* attribute.

29781 The *posix\_trace\_attr\_getmaxdatasize()* function shall copy the maximum user trace event data  
 29782 size, in bytes, from the *max-data-size* attribute of the attributes object pointed to by the *attr*  
 29783 argument into the variable pointed to by the *maxdatasize* argument. The default value for the  
 29784 *max-data-size* attribute is implementation-defined.

29785 The *posix\_trace\_attr\_getmaxsystemeventsize()* function calculates the maximum memory size, in  
 29786 bytes, required to store a single system trace event. This value is calculated for the trace stream  
 29787 attributes object pointed to by the *attr* argument and is returned in the variable pointed to by the  
 29788 *eventsize* argument.

29789 The values returned as the maximum memory sizes of the user and system trace events shall be  
 29790 such that if the sum of the maximum memory sizes of a set of the trace events that may be  
 29791 recorded in a trace stream is less than or equal to the *stream-min-size* attribute of that trace

29792 stream, the system provides the necessary resources for recording all those trace events, without  
29793 loss.

29794 The *posix\_trace\_attr\_getmaxusersize()* function calculates the maximum memory size, in  
29795 bytes, required to store a single user trace event generated by a call to *posix\_trace\_event()* with a  
29796 *data\_len* parameter equal to the *data\_len* value specified in this call. This value is calculated for  
29797 the trace stream attributes object pointed to by the *attr* argument and is returned in the variable  
29798 pointed to by the *eventsz* argument.

29799 The *posix\_trace\_attr\_getstreamsize()* function shall copy the stream size, in bytes, from the  
29800 *stream-min-size* attribute of the attributes object pointed to by the *attr* argument into the variable  
29801 pointed to by the *streamsize* argument.

29802 This stream size is the current total memory size reserved for system and user trace events in the  
29803 trace stream. The default value for the *stream-min-size* attribute is implementation-defined. The  
29804 stream size refers to memory used to store trace event records. Other stream data (for example,  
29805 trace attribute values) shall not be included in this size.

29806 The *posix\_trace\_attr\_setmaxdatasize()* function shall set the maximum allowed size, in bytes, in  
29807 the *max-data-size* attribute of the attributes object pointed to by the *attr* argument, using the size  
29808 value supplied by the *maxdatasize* argument. This maximum size is the maximum allowed size  
29809 for the user data argument which may be passed to *posix\_trace\_event()*. The implementation  
29810 shall be allowed to truncate data passed to *trace\_user\_event* which is longer than *maxdatasize*.

29811 The *posix\_trace\_attr\_setstreamsize()* function shall set the minimum allowed size, in bytes, in the  
29812 *stream-min-size* attribute of the attributes object pointed to by the *attr* argument, using the size  
29813 value supplied by the *streamsize* argument.

#### 29814 RETURN VALUE

29815 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
29816 return the corresponding error number.

29817 TRL The *posix\_trace\_attr\_getlogsize()* function stores the maximum trace log allowed size in the object  
29818 pointed to by *logsize*, if successful.

29819 The *posix\_trace\_attr\_getmaxdatasize()* function stores the maximum trace event record memory  
29820 size in the object pointed to by *maxdatasize*, if successful.

29821 The *posix\_trace\_attr\_getmaxsystemeventsz()* function stores the maximum memory size to store  
29822 a single system trace event in the object pointed to by *eventsz*, if successful.

29823 The *posix\_trace\_attr\_getmaxusersize()* function stores the maximum memory size to store a  
29824 single user trace event in the object pointed to by *eventsz*, if successful.

29825 The *posix\_trace\_attr\_getstreamsize()* function stores the maximum trace stream allowed size in  
29826 the object pointed to by *streamsize*, if successful.

#### 29827 ERRORS

29828 These functions may fail if:

29829 [EINVAL] The value specified by one of the arguments is invalid.

29830 **EXAMPLES**

29831           None.

29832 **APPLICATION USAGE**

29833           None.

29834 **RATIONALE**

29835           None.

29836 **FUTURE DIRECTIONS**

29837           None.

29838 **SEE ALSO**

29839           *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_event()*, *posix\_trace\_get\_attr()*, the Base

29840           Definitions volume of IEEE Std. 1003.1-200x, `<sys/types.h>`, `<trace.h>`

29841 **CHANGE HISTORY**

29842           First released in Issue 6. Derived from the IEEE Std. 1003.1q-2000.

29843 **NAME**

29844        posix\_trace\_attr\_init — trace stream attributes object initialization

29845 **SYNOPSIS**

29846 TRC     #include &lt;trace.h&gt;

29847        int posix\_trace\_attr\_init(trace\_attr\_t \*attr);

29848

29849 **DESCRIPTION**29850        Refer to *posix\_trace\_attr\_destroy()*.

29851 **NAME**

29852        posix\_trace\_clear — clear trace stream and trace log

29853 **SYNOPSIS**

29854 TRC     #include <sys/types.h>

29855        #include <trace.h>

29856        int posix\_trace\_clear(trace\_id\_t trid);

29857

29858 **DESCRIPTION**

29859        The *posix\_trace\_clear()* function shall reinitialize the trace stream identified by the argument *trid* as if it were returning from the *posix\_trace\_create()* function, except that the same allocated resources are reused, the mapping of trace event type identifiers to trace event names is unchanged, and the trace stream status remains unchanged (that is, if it was running, it remains running and if it was suspended, it remains suspended).

29864        All trace events in the trace stream recorded before the call to *posix\_trace\_clear()* are lost. The *posix\_stream\_full\_status* status shall be set to POSIX\_TRACE\_NOT\_FULL. There is no guarantee that all trace events that occurred during the *posix\_trace\_clear()* call are recorded; the behavior with respect to trace points that may occur during this call, is unspecified.

29868 TRL     If the Trace Log option is supported and the trace stream has been created with a log, the *posix\_trace\_clear()* function shall reinitialize the trace stream with the same behavior as if the trace stream was created without the log, plus it shall reinitialize the trace log associated with the trace stream identified by the argument *trid* as if it were returning from the *posix\_trace\_create\_withlog()* function, except that the same allocated resources, for the trace log, may be reused and the associated trace stream status remains unchanged. The first trace event recorded in the trace log after the call to *posix\_trace\_clear()* shall be the same as the first trace event recorded in the active trace stream after the call to *posix\_trace\_clear()*. The *posix\_log\_full\_status* status shall be set to POSIX\_TRACE\_NOT\_FULL. There is no guarantee that all trace events that occurred during the *posix\_trace\_clear()* call are recorded in the trace log; the behavior with respect to trace points that may occur during this call is unspecified. If the log full policy is POSIX\_TRACE\_APPEND, the effect of a call to this function is unspecified for the trace log associated with the trace stream identified by the *trid* argument.

29881 **RETURN VALUE**

29882        Upon successful completion, the *posix\_trace\_clear()* function shall return a value of zero.  
29883        Otherwise, it shall return the corresponding error number.

29884 **ERRORS**

29885        The *posix\_trace\_clear()* function may fail if:

29886        [EINVAL]        The value of the *trid* argument does not correspond to an active trace stream.

29887 **EXAMPLES**

29888        None.

29889 **APPLICATION USAGE**

29890        None.

29891 **RATIONALE**

29892        None.

29893 **FUTURE DIRECTIONS**

29894        None.

29895 **SEE ALSO**

29896        *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_flush()*, *posix\_trace\_get\_attr()*, the Base  
29897        Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <trace.h>

29898 **CHANGE HISTORY**

29899        First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

29900 **NAME**

29901 posix\_trace\_close, posix\_trace\_open, posix\_trace\_rewind — trace log management

29902 **SYNOPSIS**

29903 TRC TRL #include <trace.h>

```
29904 int posix_trace_close(trace_id_t trid);
29905 int posix_trace_open(int file_desc, trace_id_t *trid);
29906 int posix_trace_rewind(trace_id_t trid);
29907
```

29908 **DESCRIPTION**

29909 The *posix\_trace\_close()* function shall deallocate the trace log identifier indicated by *trid*, and all  
 29910 of its associated resources. If there is no valid trace log pointed to by the *trid*, this function shall  
 29911 fail.

29912 The *posix\_trace\_open()* function allocates the necessary resources and establishes the connection  
 29913 between a trace log identified by the *file\_desc* argument and a trace stream identifier identified by  
 29914 the object pointed to by the *trid* argument. The *file\_desc* argument should be a valid open file  
 29915 descriptor that corresponds to a trace log. The *file\_desc* argument shall be open for reading. The  
 29916 current trace event timestamp, which specifies the timestamp of the trace event that will be read  
 29917 by the next call to *posix\_trace\_getnext\_event()*, shall be set to the timestamp of the oldest trace  
 29918 event recorded in the trace log identified by *trid*.

29919 The *posix\_trace\_open()* function returns a trace stream identifier in the variable pointed to by the  
 29920 *trid* argument, that may only be used by the following functions:

|       |                                               |  |                                    |  |
|-------|-----------------------------------------------|--|------------------------------------|--|
| 29921 | <i>posix_trace_close()</i>                    |  | <i>posix_trace_get_attr()</i>      |  |
| 29922 | <i>posix_trace_eventid_equal()</i>            |  | <i>posix_trace_get_status()</i>    |  |
| 29923 | <i>posix_trace_eventid_get_name()</i>         |  | <i>posix_trace_getnext_event()</i> |  |
| 29924 | <i>posix_trace_eventtypelist_getnext_id()</i> |  | <i>posix_trace_rewind()</i>        |  |
| 29925 | <i>posix_trace_eventtypelist_rewind()</i>     |  |                                    |  |

29926 In particular, notice that the operations normally used by a trace controller process, such as  
 29927 *posix\_trace\_start()*, *posix\_trace\_stop()*, or *posix\_trace\_shutdown()*, cannot be invoked using the  
 29928 trace stream identifier returned by the *posix\_trace\_open()* function.

29929 The *posix\_trace\_rewind()* function shall reset the current trace event timestamp, which specifies  
 29930 the timestamp of the trace event that will be read by the next call to *posix\_trace\_getnext\_event()*,  
 29931 to the timestamp of the oldest trace event recorded in the trace log identified by *trid*.

29932 **RETURN VALUE**

29933 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 29934 return the corresponding error number.

29935 If successful, the *posix\_trace\_open()* function stores the trace stream identifier value in the object  
 29936 pointed to by *trid*.

29937 **ERRORS**

- 29938 The *posix\_trace\_open()* function may fail if:
- 29939 [EINTR] The operation was interrupted by a signal and thus no trace log was opened.
  - 29940 [EINVAL] The object pointed to by *file\_desc* does not correspond to a valid trace log.
- 29941 The *posix\_trace\_close()* and *posix\_trace\_rewind()* functions may fail if:
- 29942 [EINVAL] The object pointed to by *trid* does not correspond to a valid trace log.

29943 **EXAMPLES**

29944 None.

29945 **APPLICATION USAGE**

29946 None.

29947 **RATIONALE**

29948 None.

29949 **FUTURE DIRECTIONS**

29950 None.

29951 **SEE ALSO**

29952 *posix\_trace\_get\_attr()*, *posix\_trace\_get\_filter()*, *posix\_trace\_getnext\_event()*, the Base Definitions  
29953 volume of IEEE Std. 1003.1-200x, <trace.h>

29954 **CHANGE HISTORY**

29955 First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

29956 NAME

29957 posix\_trace\_create, posix\_trace\_create\_withlog, posix\_trace\_flush, posix\_trace\_shutdown —  
 29958 trace stream initialization, flush, and shutdown from a process

29959 SYNOPSIS

```

29960 TRC #include <sys/types.h>
29961 #include <trace.h>

29962 int posix_trace_create(pid_t pid, const trace_attr_t *attr,
29963 trace_id_t *trid);
29964 TRC TRL int posix_trace_create_withlog(pid_t pid, const trace_attr_t *attr,
29965 int file_desc, trace_id_t *trid);
29966 int posix_trace_flush(trace_id_t trid);
29967 int posix_trace_shutdown(trace_id_t trid);
29968

```

29969 DESCRIPTION

29970 The *posix\_trace\_create()* function creates an active trace stream. It allocates all the resources  
 29971 needed by the trace stream being created for tracing the process specified by *pid* in accordance  
 29972 with the *attr* argument. The *attr* argument represents the initial attributes of the trace stream and  
 29973 shall have been initialized by the function *posix\_trace\_attr\_init()* prior the *posix\_trace\_create()*  
 29974 call. If the argument *attr* is NULL, the default attributes shall be used. The *attr* attributes object  
 29975 shall be manipulated through a set of functions described in the *posix\_trace\_attr* family of  
 29976 functions. If the attributes of the object pointed to by *attr* are modified later, the attributes of the  
 29977 trace stream shall not be affected. The *creation-time* attribute of the newly created trace stream  
 29978 shall be set to the value of the system clock, if the Timers option is not supported, or to the value  
 29979 of the CLOCK\_REALTIME clock, if the Timers option is supported.

29980 The *pid* argument represents the target process to be traced. If the process executing this  
 29981 function does not have appropriate privileges to trace the process identified by *pid*, an error shall  
 29982 be returned. If the *pid* argument is zero, the calling process shall be traced.

29983 The *posix\_trace\_create()* function stores the trace stream identifier of the new trace stream in the  
 29984 object pointed to by the *trid* argument. This trace stream identifier shall be used in subsequent  
 29985 calls to control tracing. The *trid* argument may only be used by the following functions:

|       |                                               |  |                                         |  |
|-------|-----------------------------------------------|--|-----------------------------------------|--|
| 29986 | <i>posix_trace_clear()</i>                    |  | <i>posix_trace_getnext_event()</i>      |  |
| 29987 | <i>posix_trace_eventid_equal()</i>            |  | <i>posix_trace_shutdown()</i>           |  |
| 29988 | <i>posix_trace_eventid_get_name()</i>         |  | <i>posix_trace_start()</i>              |  |
| 29989 | <i>posix_trace_eventtypelist_getnext_id()</i> |  | <i>posix_trace_stop()</i>               |  |
| 29990 | <i>posix_trace_eventtypelist_rewind()</i>     |  | <i>posix_trace_timedgetnext_event()</i> |  |
| 29991 | <i>posix_trace_get_attr()</i>                 |  | <i>posix_trace_trid_eventid_open()</i>  |  |
| 29992 | <i>posix_trace_get_status()</i>               |  | <i>posix_trace_trygetnext_event()</i>   |  |

29993 TEF If the Trace Event Filter option is supported, the following additional functions may use the *trid*  
 29994 argument:

|       |                                 |  |                                 |  |
|-------|---------------------------------|--|---------------------------------|--|
| 29995 | <i>posix_trace_get_filter()</i> |  | <i>posix_trace_spt_filter()</i> |  |
|-------|---------------------------------|--|---------------------------------|--|

29996

29997 In particular, notice that the operations normally used by a trace analyser process, such as  
 29998 *posix\_trace\_rewind()* or *posix\_trace\_close()*, cannot be invoked using the trace stream identifier  
 29999 returned by the *posix\_trace\_create()* function.

- 30000 TEF A trace stream shall be created in a suspended state. If the Trace Event Filter option is supported, its trace event type filter shall be empty.
- 30001
- 30002 The *posix\_trace\_create()* function may be called multiple times from the same or different processes, with the system-wide limit indicated by the runtime invariant value {TRACE\_SYS\_MAX}, which has the minimum value {\_POSIX\_TRACE\_SYS\_MAX}.
- 30003
- 30004
- 30005 The trace stream identifier returned by the *posix\_trace\_create()* function in the argument pointed to by *trid* is valid only in the process that made the function call. If it is used from another process, that is a child process, in functions defined in IEEE Std. 1003.1-200x, these functions shall return with the error [EINVAL].
- 30006
- 30007
- 30008
- 30009 TRL The *posix\_trace\_create\_withlog()* function creates a trace stream as in the *posix\_trace\_create()* function and behaves the same way, plus it associates a trace log with this trace stream. The *file\_desc* argument shall be the file descriptor designating the trace log destination. The function shall fail if this file descriptor refers to a file with a file type that is not compatible with the log policy associated with the trace log. The list of the appropriate file types that are compatible with each log policy shall be implementation-defined.
- 30010
- 30011
- 30012
- 30013
- 30014
- 30015 The *posix\_trace\_create\_withlog()* function returns in the parameter pointed to by *trid* the trace stream identifier, which uniquely identifies the newly created trace stream, and shall be used in subsequent calls to control tracing. The *trid* argument may only be used by the following functions:
- 30016
- 30017
- 30018
- |       |                                               |                                         |  |
|-------|-----------------------------------------------|-----------------------------------------|--|
| 30019 | <i>posix_trace_clear()</i>                    | <i>posix_trace_getnext_event()</i>      |  |
| 30020 | <i>posix_trace_eventid_equal()</i>            | <i>posix_trace_shutdown()</i>           |  |
| 30021 | <i>posix_trace_eventid_get_name()</i>         | <i>posix_trace_start()</i>              |  |
| 30022 | <i>posix_trace_eventtypelist_getnext_id()</i> | <i>posix_trace_stop()</i>               |  |
| 30023 | <i>posix_trace_eventtypelist_rewind()</i>     | <i>posix_trace_timedgetnext_event()</i> |  |
| 30024 | <i>posix_trace_flush()</i>                    | <i>posix_trace_trid_eventid_open()</i>  |  |
| 30025 | <i>posix_trace_get_attr()</i>                 | <i>posix_trace_trygetnext_event()</i>   |  |
| 30026 | <i>posix_trace_get_status()</i>               |                                         |  |
- 30027
- 30028 TRL TEF If the Trace Event Filter option is supported, the following additional functions may use the *trid* argument:
- 30029
- |       |                                 |                                 |  |
|-------|---------------------------------|---------------------------------|--|
| 30030 | <i>posix_trace_get_filter()</i> | <i>posix_trace_spt_filter()</i> |  |
|-------|---------------------------------|---------------------------------|--|
- 30031 In particular, notice that the operations normally used by a trace analyser process, such as *posix\_trace\_rewind()* or *posix\_trace\_close()*, cannot be invoked using the trace stream identifier returned by the *posix\_trace\_create\_withlog()* function.
- 30032
- 30033
- 30034 The *posix\_trace\_flush()* function initiates a flush operation which copies the contents of the trace stream identified by the argument *trid* into the trace log associated with the trace stream at the creation time. If no trace log has been associated with the trace stream pointed to by *trid*, this function shall return an error. The termination of the flush operation can be polled by the *posix\_trace\_get\_status()* function. During the flush operation, it shall be possible to trace new trace events up to the point when the trace stream becomes full. After flushing is completed, the space used by the flushed trace events shall be available for tracing new trace events.
- 30035
- 30036
- 30037
- 30038
- 30039
- 30040
- 30041 If flushing the trace stream causes the resulting trace log to become full, the trace log full policy shall be applied. If the trace *log-full-policy* attribute is set, the following occurs:
- 30042

|       |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 30043 | POSIX_TRACE_UNTIL_FULL |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30044 |                        | The trace events that have not yet been flushed are discarded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 30045 | POSIX_TRACE_LOOP       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30046 |                        | The trace events that have not yet been flushed are written to the beginning of the trace log, overwriting previous trace events stored there.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 30047 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30048 | POSIX_TRACE_APPEND     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30049 |                        | The trace events that had not yet been flushed shall be appended to the trace log.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| 30050 |                        | For an active trace stream with log, when the <i>posix_trace_shutdown()</i> function is called, all trace events that have not yet been flushed to the trace log shall be flushed, as in the <i>posix_trace_flush()</i> function, and the trace log shall be closed.                                                                                                                                                                                                                                                                                                                              |
| 30051 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30052 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30053 |                        | When a trace log is closed, all the information that may be retrieved later from the trace log through the trace interface, shall have been written to the trace log. This information includes the trace attributes, the list of trace event types (with the mapping between trace event names and trace event type identifiers), and the trace status.                                                                                                                                                                                                                                          |
| 30054 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30055 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30056 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30057 |                        | In addition, some unspecified information shall be written to the trace log to allow detection of a valid trace log during the <i>posix_trace_open()</i> operation.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 30058 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30059 |                        | The <i>posix_trace_shutdown()</i> function shall stop the tracing of trace events in the trace stream identified by <i>trid</i> , as if <i>posix_trace_stop()</i> had been invoked. The <i>posix_trace_shutdown()</i> function shall free all the resources associated with the trace stream.                                                                                                                                                                                                                                                                                                     |
| 30060 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30061 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30062 |                        | The <i>posix_trace_shutdown()</i> function shall not return until all the resources associated with the trace stream have been freed. When the <i>posix_trace_shutdown()</i> function returns, the <i>trid</i> argument becomes an invalid trace stream identifier. A call to this function shall unconditionally deallocate the resources regardless of whether all trace events have been retrieved by the analyzer process. Any thread blocked on one of the <i>trace_getnext_event()</i> functions (which specified this <i>trid</i> ) before this call is unblocked with the error [EINVAL]. |
| 30063 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30064 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30065 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30066 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30067 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30068 |                        | If the process exits, invokes an <i>exec()</i> call, or is terminated, the trace streams that the process had created and that have not yet been shut down, shall be automatically shut down as if an explicit call were made to the <i>posix_trace_shutdown()</i> function.                                                                                                                                                                                                                                                                                                                      |
| 30069 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30070 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30071 |                        | The <i>posix_trace_shutdown()</i> function shall not return until all trace events have been flushed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 30072 | <b>RETURN VALUE</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30073 |                        | Upon successful completion, these functions shall return a value of zero. Otherwise, they shall return the corresponding error number.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 30074 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30075 | TRL                    | The <i>posix_trace_create()</i> and <i>posix_trace_create_withlog()</i> functions store the trace stream identifier value in the object pointed to by <i>trid</i> , if successful.                                                                                                                                                                                                                                                                                                                                                                                                                |
| 30076 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30077 | <b>ERRORS</b>          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30078 | TRL                    | The <i>posix_trace_create()</i> and <i>posix_trace_create_withlog()</i> functions shall fail if:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| 30079 | [EAGAIN]               | No more trace streams can be started now. {TRACE_SYS_MAX} has been exceeded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 30080 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30081 | [EINTR]                | The operation was interrupted by a signal. No trace stream was created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 30082 | [EINVAL]               | One or more of the trace parameters specified by the <i>attr</i> parameter is invalid.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 30083 | [ENOMEM]               | The implementation does not currently have sufficient memory to create the trace stream with the specified parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 30084 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 30085 | [EPERM]                | The caller does not have appropriate privilege to trace the process specified by <i>pid</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 30086 |                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

|       |          |                                                                                                                                                  |
|-------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| 30087 | [ESRCH]  | The <i>pid</i> argument does not refer to an existing process.                                                                                   |
| 30088 | TRL      | The <i>posix_trace_create_withlog()</i> function shall fail if:                                                                                  |
| 30089 | [EBADF]  | The <i>file_desc</i> argument is not a valid file descriptor open for writing.                                                                   |
| 30090 | [EINVAL] | The <i>file_desc</i> argument refers to a file with a file type that does not support the log policy associated with the trace log.              |
| 30091 |          |                                                                                                                                                  |
| 30092 | [ENOSPC] | No space left on device. The device corresponding to the argument <i>file_desc</i> does not contain the space required to create this trace log. |
| 30093 |          |                                                                                                                                                  |
| 30094 |          |                                                                                                                                                  |
| 30095 | TRL      | The <i>posix_trace_flush()</i> and <i>posix_trace_shutdown()</i> functions shall fail if:                                                        |
| 30096 | [EINVAL] | The value of the <i>trid</i> argument does not correspond to an active trace stream with log.                                                    |
| 30097 |          |                                                                                                                                                  |
| 30098 | [EFBIG]  | The trace log file has attempted to exceed an implementation-defined maximum file size.                                                          |
| 30099 |          |                                                                                                                                                  |
| 30100 | [ENOSPC] | No space left on device.                                                                                                                         |
| 30101 |          |                                                                                                                                                  |
| 30102 |          | <b>EXAMPLES</b>                                                                                                                                  |
| 30103 |          | None.                                                                                                                                            |
| 30104 |          | <b>APPLICATION USAGE</b>                                                                                                                         |
| 30105 |          | None.                                                                                                                                            |
| 30106 |          | <b>RATIONALE</b>                                                                                                                                 |
| 30107 |          | None.                                                                                                                                            |
| 30108 |          | <b>FUTURE DIRECTIONS</b>                                                                                                                         |
| 30109 |          | None.                                                                                                                                            |
| 30110 |          | <b>SEE ALSO</b>                                                                                                                                  |
| 30111 |          | <i>clock_getres()</i> , <i>exec</i> , <i>posix_trace_attr_init()</i> , <i>posix_trace_clear()</i> , <i>posix_trace_close()</i> ,                 |
| 30112 |          | <i>posix_trace_eventid_equal()</i> , <i>posix_trace_eventtypelist_getnext_id()</i> , <i>posix_trace_flush()</i> ,                                |
| 30113 |          | <i>posix_trace_get_attr()</i> , <i>posix_trace_get_filter()</i> , <i>posix_trace_get_status()</i> , <i>posix_trace_getnext_event()</i> ,         |
| 30114 |          | <i>posix_trace_open()</i> , <i>posix_trace_rewind()</i> , <i>posix_trace_set_filter()</i> , <i>posix_trace_shutdown()</i> ,                      |
| 30115 |          | <i>posix_trace_start()</i> , <i>posix_trace_timedgetnext_event()</i> , <i>posix_trace_trid_eventid_open()</i> ,                                  |
| 30116 |          | <i>posix_trace_start()</i> , <i>time()</i> , the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>,                                |
| 30117 |          | <trace.h>                                                                                                                                        |
| 30118 |          | <b>CHANGE HISTORY</b>                                                                                                                            |
| 30119 |          | First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.                                                                                  |

30120 **NAME**

30121 posix\_trace\_event, posix\_trace\_eventid\_open — trace functions for instrumenting application  
 30122 code

30123 **SYNOPSIS**

```
30124 TRC #include <sys/types.h>
30125 #include <trace.h>

30126 void posix_trace_event(trace_event_id_t event_id, const void *data_ptr,
30127 size_t data_len);
30128 int posix_trace_eventid_open(const char *event_name,
30129 trace_event_id_t *event_id);
30130
```

30131 **DESCRIPTION**

30132 The *posix\_trace\_event()* function records the *event\_id* and the user data pointed to by *data\_ptr* in  
 30133 the trace stream into which the calling process is being traced and in which *event\_id* is not  
 30134 filtered out. If the total size of the user trace event data represented by *data\_len* is not greater  
 30135 than the declared maximum size for user trace event data, then the *truncation-status* attribute of  
 30136 the trace event recorded is POSIX\_TRACE\_NOT\_TRUNCATED. Otherwise, the user trace event  
 30137 data is truncated to this declared maximum size and the *truncation-status* attribute of the trace  
 30138 event recorded is POSIX\_TRACE\_TRUNCATED\_RECORD.

30139 If there is no trace stream created for the process or if the created trace stream is not running or if  
 30140 the trace event specified by *event\_id* is filtered out in the trace stream, the *posix\_trace\_event()*  
 30141 function has no effect.

30142 The *posix\_trace\_eventid\_open()* function is used to associate a user trace event name with a trace  
 30143 event type identifier for the calling process. The trace event name is the string pointed to by the  
 30144 argument *event\_name*. It shall have a maximum of {TRACE\_EVENT\_NAME\_MAX} characters  
 30145 (which has the minimum value {\_POSIX\_TRACE\_EVENT\_NAME\_MAX}). The number of user  
 30146 trace event type identifiers that can be defined for any given process is limited by the maximum  
 30147 value {TRACE\_USER\_EVENT\_MAX}, which has the minimum value  
 30148 {POSIX\_TRACE\_USER\_EVENT\_MAX}.

30149 If the Trace Inheritance option is not supported, the *posix\_trace\_trid\_eventid\_open()* function shall  
 30150 associate the user trace event name pointed to by the *event\_name* argument with a trace event  
 30151 type identifier that is unique for the traced process, and is returned in the variable pointed to by  
 30152 the *event* argument. If the user trace event name has already been mapped for the traced process,  
 30153 then the previously assigned trace event type identifier shall be returned. If the per-process user  
 30154 trace event name limit represented by {TRACE\_USER\_EVENT\_MAX} has been reached, the  
 30155 pre-defined POSIX\_TRACE\_UNNAMED\_USEREVENT (see Table 2-10 (on page 582)) user trace  
 30156 event shall be returned.

30157 TRI If the Trace Inherit option is supported, the *posix\_trace\_trid\_eventid\_open()* function shall  
 30158 associate the user trace event name pointed to by the *event\_name* argument with a trace event  
 30159 type identifier that is unique for all the processes being traced in this same trace stream, and is  
 30160 returned in the variable pointed to by the *event* argument. If the user trace event name has  
 30161 already been mapped for the traced processes, then the previously assigned trace event type  
 30162 identifier shall be returned. If the per-process user trace event name limit represented by  
 30163 {TRACE\_USER\_EVENT\_MAX} has been reached, the pre-defined  
 30164 POSIX\_TRACE\_UNNAMED\_USEREVENT (Table 2-10 (on page 582)) user trace event shall be  
 30165 returned.

30166 **Note:** The above procedure, together with the fact that multiple processes can only be  
 30167 traced into the same trace stream by inheritance, ensure that all the processes that are

30168                   traced into a trace stream have the same mapping of trace event names to trace event  
30169                   type identifiers.

30170

30171                   If there is no trace stream created, the *posix\_trace\_eventid\_open()* function shall store this  
30172                   information for future trace streams created for this process.

30173 **RETURN VALUE**

30174                   No return value is defined for the *posix\_trace\_event()* function.

30175                   Upon successful completion, the *posix\_trace\_eventid\_open()* function shall return a value of zero.  
30176                   Otherwise, it shall return the corresponding error number. The *posix\_trace\_eventid\_open()*  
30177                   function stores the trace event type identifier value in the object pointed to by *event\_id*, if  
30178                   successful.

30179 **ERRORS**

30180                   The *posix\_trace\_eventid\_open()* function may fail if:

30181                   [ENAMETOOLONG]

30182                                           The size of the name pointed to by *event\_name* argument was longer than the  
30183                                           implementation-defined value {TRACE\_EVENT\_NAME\_MAX}.

30184 **EXAMPLES**

30185                   None.

30186 **APPLICATION USAGE**

30187                   None.

30188 **RATIONALE**

30189                   None.

30190 **FUTURE DIRECTIONS**

30191                   None.

30192 **SEE ALSO**

30193                   *posix\_trace\_start()*, *posix\_trace\_trid\_eventid\_open()*, the Base Definitions volume of  
30194                   IEEE Std. 1003.1-200x, <sys/types.h>, <trace.h>

30195 **CHANGE HISTORY**

30196                   First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

30197 NAME

30198 posix\_trace\_eventid\_equal, posix\_trace\_eventid\_get\_name, posix\_trace\_trid\_eventid\_open —  
 30199 manipulate trace event type identifier

30200 SYNOPSIS

```
30201 TRC #include <trace.h>

30202 int posix_trace_eventid_equal(trace_id_t trid, trace_eventid_t event1,
30203 trace_eventid_t event2);
30204 int posix_trace_eventid_get_name(trace_id_t trid, trace_eventid_t event,
30205 char *event_name);
30206 TRC TEF int posix_trace_trid_eventid_open(trace_id_t trid,
30207 const char *event_name, trace_eventid_t *event);
30208
```

30209 DESCRIPTION

30210 The *posix\_trace\_eventid\_equal()* function compares the trace event type identifiers *event1* and  
 30211 *event2* from the same trace stream or the same trace log identified by the *trid* argument. If the  
 30212 trace event type identifiers *event1* and *event2* are from different trace streams, the return value  
 30213 shall be unspecified.

30214 The *posix\_trace\_eventid\_get\_name()* function returns in the argument pointed to by *event\_name*,  
 30215 the trace event name associated with the trace event type identifier identified by the argument  
 30216 *event*, for the trace stream or for the trace log identified by the *trid* argument. The name of the  
 30217 trace event shall have a maximum of {TRACE\_EVENT\_NAME\_MAX} characters (which has the  
 30218 minimum value {\_POSIX\_TRACE\_EVENT\_NAME\_MAX}). Successive calls to this function  
 30219 with the same trace event type identifier and the same trace stream identifier shall return the  
 30220 same event name.

30221 TEF The *posix\_trace\_trid\_eventid\_open()* function is used to associate a user trace event name with a  
 30222 trace event type identifier for a given trace stream. The trace stream is identified by the *trid*  
 30223 argument, and it shall be an active trace stream. The trace event name is the string pointed to by  
 30224 the argument *event\_name*. It shall have a maximum of {TRACE\_EVENT\_NAME\_MAX}  
 30225 characters (which has the minimum value {\_POSIX\_TRACE\_EVENT\_NAME\_MAX}). The  
 30226 number of user trace event type identifiers that can be defined for any given process is limited  
 30227 by the maximum value {TRACE\_USER\_EVENT\_MAX}, which has the minimum value  
 30228 {\_POSIX\_TRACE\_USER\_EVENT\_MAX}.

30229 If the Trace Inheritance option is not supported, the *posix\_trace\_trid\_eventid\_open()* function shall  
 30230 associate the user trace event name pointed to by the *event\_name* argument with a trace event  
 30231 type identifier that is unique for the process being traced in the trace stream identified by the *trid*  
 30232 argument, and is returned in the variable pointed to by the *event* argument. If the user trace  
 30233 event name has already been mapped for the traced process, then the previously assigned trace  
 30234 event type identifier shall be returned. If the per-process user trace event name limit represented  
 30235 by {TRACE\_USER\_EVENT\_MAX} has been reached, the pre-defined  
 30236 POSIX\_TRACE\_UNNAMED\_USEREVENT (see Table 2-10 (on page 582)) user trace event shall  
 30237 be returned.

30238 TEF TRI If the Trace Inheritance option is supported, the *posix\_trace\_trid\_eventid\_open()* function shall  
 30239 associate the user trace event name pointed to by the *event\_name* argument with a trace event  
 30240 type identifier that is unique for all the processes being traced in the trace stream identified by  
 30241 the *trid* argument, and is returned in the variable pointed to by the *event* argument. If the user  
 30242 trace event name has already been mapped for the traced processes, then the previously  
 30243 assigned trace event type identifier shall be returned. If the per-process user trace event name  
 30244 limit represented by {TRACE\_USER\_EVENT\_MAX} has been reached, the pre-defined  
 30245 POSIX\_TRACE\_UNNAMED\_USEREVENT (see Table 2-10 (on page 582)) user trace event shall

30246 be returned.

### 30247 RETURN VALUE

30248 TEF Upon successful completion, the `posix_trace_eventid_get_name()` and  
30249 `posix_trace_trid_eventid_open()` functions shall return a value of zero. Otherwise, they shall return  
30250 the corresponding error number.

30251 The `posix_trace_eventid_equal()` function shall return a non-zero value if *event1* and *event2* are  
30252 equal; otherwise, a value of zero shall be returned. No errors are defined. If either *event1* or  
30253 *event2* are not valid trace event type identifiers for the trace stream specified by *trid* or if the *trid*  
30254 is invalid, the behavior shall be unspecified.

30255 The `posix_trace_eventid_get_name()` function stores the trace event name value in the object  
30256 pointed to by *event\_name*, if successful.

30257 TEF The `posix_trace_trid_eventid_open()` function stores the trace event type identifier value in the  
30258 object pointed to by *event*, if successful.

### 30259 ERRORS

30260 TEF The `posix_trace_eventid_get_name()` and `posix_trace_trid_eventid_open()` functions may fail if:

30261 [EINVAL] The *trid* argument was not a valid trace type identifier.

30262 TEF The `posix_trace_trid_eventid_open()` function may fail if:

30263 [ENAMETOOLONG]

30264 The size of the name pointed to by *event\_name* argument was longer than the  
30265 implementation-defined value {TRACE\_EVENT\_NAME\_MAX}.  
30266

30267 The `posix_trace_eventid_get_name()` function may fail if:

30268 [EINVAL] The trace event type identifier *event* was not associated with any name.

### 30269 EXAMPLES

30270 None.

### 30271 APPLICATION USAGE

30272 None.

### 30273 RATIONALE

30274 None.

### 30275 FUTURE DIRECTIONS

30276 None.

### 30277 SEE ALSO

30278 `posix_trace_event()`, <REFERENCE UNDEFINED>(posix\_trace\_eventid),  
30279 `posix_trace_getnext_event()`, the Base Definitions volume of IEEE Std. 1003.1-200x, <trace.h>

### 30280 CHANGE HISTORY

30281 First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

## 30282 NAME

30283 `posix_trace_eventset_add`, `posix_trace_eventset_del`, `posix_trace_eventset_empty`,  
 30284 `posix_trace_eventset_fill`, `posix_trace_eventset_ismember` — manipulate trace event type sets

## 30285 SYNOPSIS

```
30286 TRC TEF #include <trace.h>
30287
30287 int posix_trace_eventset_add(trace_event_id_t event_id,
30288 trace_event_set_t *set);
30289 int posix_trace_eventset_del(trace_event_id_t event_id,
30290 trace_event_set_t *set);
30291 int posix_trace_eventset_empty(trace_event_set_t *set);
30292 int posix_trace_eventset_fill(trace_event_set_t *set, int what);
30293 int posix_trace_eventset_ismember(trace_event_id_t event_id,
30294 const trace_event_set_t *set, int *ismember);
30295
```

## 30296 DESCRIPTION

30297 These primitives manipulate sets of trace event types. They operate on data objects addressable  
 30298 by the application, not on the current trace event filter of any trace stream.

30299 The `posix_trace_eventset_add()` and `posix_trace_eventset_del()` functions, respectively, add or  
 30300 delete the individual trace event type specified by the value of the argument `event_id` to or from  
 30301 the trace event type set pointed to by the argument `set`. Adding a trace event type already in the  
 30302 set or deleting a trace event type not in the set shall not be considered an error.

30303 The `posix_trace_eventset_empty()` function initializes the trace event type set pointed to by the `set`  
 30304 argument such that all trace event types defined, both system and user, shall be excluded from  
 30305 the set.

30306 The `posix_trace_eventset_fill()` function initializes the trace event type set pointed to by the  
 30307 argument `set`, such that the set of trace event types defined by the argument `what` shall be  
 30308 included in the set. The value of the argument `what` shall consist of one of the following values,  
 30309 as defined in the `<trace.h>` header:

30310 **POSIX\_TRACE\_WOPID\_EVENTS**

30311 All the process-independent implementation-defined system trace event types are included  
 30312 in the set.

30313 **POSIX\_TRACE\_SYSTEM\_EVENTS** All the implementation-defined system trace event types are  
 30314 included in the set, as are those defined in IEEE Std. 1003.1-200x.

30315 **POSIX\_TRACE\_ALL\_EVENTS** All trace event types defined, both system and user, are included  
 30316 in the set.

30317 Applications shall call either `posix_trace_eventset_empty()` or `posix_trace_eventset_fill()` at least  
 30318 once for each object of type `trace_event_set_t` prior to any other use of that object. If such an  
 30319 object is not initialized in this way, but is nonetheless supplied as an argument to any of the  
 30320 `posix_trace_eventset_add()`, `posix_trace_eventset_del()`, or `posix_trace_eventset_ismember()` functions,  
 30321 the results are undefined.

30322 The `posix_trace_eventset_ismember()` function tests whether the trace event type specified by the  
 30323 value of the argument `event_id` is a member of the set pointed to by the argument `set`. The value  
 30324 returned in the object pointed to by `ismember` argument is zero if the trace event type identifier is  
 30325 not a member of the set and a value different from zero if it is a member of the set.

**30326 RETURN VALUE**

30327       Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
30328       return the corresponding error number.

**30329 ERRORS**

30330       These functions may fail if:

30331       [EINVAL]       The value of one of the arguments is invalid.

**30332 EXAMPLES**

30333       None.

**30334 APPLICATION USAGE**

30335       None.

**30336 RATIONALE**

30337       None.

**30338 FUTURE DIRECTIONS**

30339       None.

**30340 SEE ALSO**

30341       *posix\_trace\_set\_filter()*, *posix\_trace\_trid\_eventid\_open()*, the Base Definitions volume of  
30342       IEEE Std. 1003.1-200x, <trace.h>

**30343 CHANGE HISTORY**

30344       First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

30345 **NAME**

30346 posix\_trace\_eventtypelist\_getnext\_id, posix\_trace\_eventtypelist\_rewind — iterate over a  
30347 mapping of trace event types

30348 **SYNOPSIS**

```
30349 TRC #include <trace.h>
30350 int posix_trace_eventtypelist_getnext_id(trace_id_t trid,
30351 trace_eventid_t *event, int *unavailable);
30352 int posix_trace_eventtypelist_rewind(trace_id_t trid);
30353
```

30354 **DESCRIPTION**

30355 The first time *posix\_trace\_eventtypelist\_getnext\_id()* is called, the function shall return in the  
30356 variable pointed to by *event* the first trace event type identifier of the list of trace events of the  
30357 trace stream identified by the *trid* argument. Successive calls to  
30358 *posix\_trace\_eventtypelist\_getnext\_id()* return in the variable pointed to by *event* the next trace  
30359 event type identifier in that same list. Each time a trace event type identifier is successfully  
30360 written into the variable pointed to by the *event* argument, the variable pointed to by the  
30361 *unavailable* argument shall be set to zero. When no more trace event type identifiers are  
30362 available, and so none is returned, the variable pointed to by the *unavailable* argument shall be  
30363 set to a value different from zero.

30364 The *posix\_trace\_eventtypelist\_rewind()* function shall reset the next trace event type identifier to  
30365 be read to the first trace event type identifier from the list of trace events used in the trace stream  
30366 identified by *trid*.

30367 **RETURN VALUE**

30368 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
30369 return the corresponding error number.

30370 The *posix\_trace\_eventtypelist\_getnext\_id()* function stores the trace event type identifier value in  
30371 the object pointed to by *event*, if successful.

30372 **ERRORS**

30373 These functions may fail if:

30374 [EINVAL] The *trid* argument was not a valid trace stream identifier.

30375 **EXAMPLES**

30376 None.

30377 **APPLICATION USAGE**

30378 None.

30379 **RATIONALE**

30380 None.

30381 **FUTURE DIRECTIONS**

30382 None.

30383 **SEE ALSO**

30384 *posix\_trace\_event()*, <REFERENCE UNDEFINED>(posix\_trace\_eventid),  
30385 *posix\_trace\_getnext\_event()*, *posix\_trace\_trid\_eventid\_open()*, the Base Definitions volume of  
30386 IEEE Std. 1003.1-200x, <trace.h>

30387 **CHANGE HISTORY**

30388 First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

30389 **NAME**

30390        posix\_trace\_flush — trace stream flush from a process

30391 **SYNOPSIS**

30392 TRC     #include <sys/types.h>

30393        #include <trace.h>

30394        int posix\_trace\_flush(trace\_id\_t trid);

30395

30396 **DESCRIPTION**

30397        Refer to *posix\_trace\_create()*.

30398 **NAME**

30399        posix\_trace\_get\_attr, posix\_trace\_get\_status — retrieve the trace attributes or trace statuses

30400 **SYNOPSIS**

30401 TRC     #include &lt;trace.h&gt;

30402        int posix\_trace\_get\_attr(trace\_id\_t trid, trace\_attr\_t \*attr);

30403        int posix\_trace\_get\_status(trace\_id\_t trid,

30404            struct posix\_trace\_status\_info \*statusinfo);

30405

30406 **DESCRIPTION**30407        The *posix\_trace\_get\_attr()* function shall copy the attributes of the active trace stream identified  
30408 TRL     by *trid* into the object pointed to by the *attr* argument. If the Trace Log option is supported, *trid*  
30409     may represent a pre-recorded trace log.30410        The *posix\_trace\_get\_status()* function returns, in the structure pointed to by the *statusinfo*  
30411     argument, the current trace status for the trace stream identified by the *trid* argument. These  
30412     status values returned in the structure pointed to by *statusinfo* shall have been appropriately  
30413 TRL     read to ensure that the returned values are consistent. If the Trace Log option is supported and  
30414     the *trid* argument refers to a pre-recorded trace stream, the status shall be the status of the  
30415     completed trace stream.30416        Each time the *posix\_trace\_get\_status()* function is used, the overrun status of the trace stream  
30417 TRL     shall be reset to POSIX\_TRACE\_NO\_OVERRUN immediately after the call completes. If the  
30418     Trace Log option is supported, the *posix\_trace\_get\_status()* function shall behave the same as  
30419     when the option is not supported except for the following differences:30420        

- If the *trid* argument refers to a trace stream with log, each time the *posix\_trace\_get\_status()*  
30421     function is used, the log overrun status of the trace stream shall be reset to  
30422     POSIX\_TRACE\_NO\_OVERRUN and the *flush\_error* status shall be reset to zero immediately  
30423     after the call completes.

30424        

- If the *trid* argument refers to a pre-recorded trace stream, the status returned shall be the  
30425     status of the completed trace stream and the status values of the trace stream shall not be  
30426     reset.

30427

30428 **RETURN VALUE**30429        Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
30430     return the corresponding error number.30431        The *posix\_trace\_get\_attr()* function stores the trace attributes in the object pointed to by *attr*, if  
30432     successful.30433        The *posix\_trace\_get\_status()* function stores the trace status in the object pointed to by *statusinfo*,  
30434     if successful.30435 **ERRORS**

30436        These functions may fail if:

30437        [EINVAL]        The trace stream argument *trid* does not correspond to a valid active trace  
30438     stream or a valid trace log.

30439 **EXAMPLES**

30440           None.

30441 **APPLICATION USAGE**

30442           None.

30443 **RATIONALE**

30444           None.

30445 **FUTURE DIRECTIONS**

30446           None.

30447 **SEE ALSO**

30448           *posix\_trace\_attr\_destroy()*, *posix\_trace\_attr\_init()*, *posix\_trace\_create()*, *posix\_trace\_open()*, the Base

30449           Definitions volume of IEEE Std. 1003.1-200x, <**trace.h**>

30450 **CHANGE HISTORY**

30451           First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

30452 **NAME**

30453 posix\_trace\_get\_filter, posix\_trace\_set\_filter — retrieve and set filter of an initialized trace  
 30454 stream

30455 **SYNOPSIS**

30456 TRC TEF #include <trace.h>

30457 int posix\_trace\_get\_filter(trace\_id\_t trid, trace\_event\_set\_t \*set);

30458 int posix\_trace\_set\_filter(trace\_id\_t trid,

30459 const trace\_event\_set\_t \*set, int how);

30460

30461 **DESCRIPTION**

30462 The *posix\_trace\_get\_filter()* function shall be used to retrieve, into the argument pointed to by *set*,  
 30463 the actual trace event filter from the trace stream specified by *trid*.

30464 The *posix\_trace\_set\_filter()* function shall be used to change the set of filtered trace event types  
 30465 after a trace stream identified by the *trid* argument is created. This function may be called prior  
 30466 to starting the trace stream, or while the trace stream is active. By default, if no call is made to  
 30467 *posix\_trace\_set\_filter()*, all trace events shall be recorded (that is, none of the trace event types are  
 30468 filtered out).

30469 If this function is called while the trace is in progress, a special system trace event,  
 30470 POSIX\_TRACE\_FILTER, shall be recorded in the trace indicating both the old and the new sets  
 30471 of filtered trace event types (see Table 2-7 (on page 580) and Table 2-9 (on page 581)).

30472 If the *posix\_trace\_set\_filter()* function is interrupted by a signal, an error is returned and the filter  
 30473 is not changed. In this case, the state of the trace stream shall not be changed.

30474 The value of the argument *how* indicates the manner in which the set is to be changed and shall  
 30475 have one of the following values, as defined in the <trace.h> header:

30476 POSIX\_TRACE\_SET\_EVENTSET

30477 The resulting set of trace event types to be filtered shall be the trace event type set pointed  
 30478 to by the argument *set*.

30479 POSIX\_TRACE\_ADD\_EVENTSET

30480 The resulting set of trace event types to be filtered shall be the union of the current set and  
 30481 the trace event type set pointed to by the argument *set*.

30482 POSIX\_TRACE\_SUB\_EVENTSET

30483 The resulting set of trace event types to be filtered shall be all trace event types in the  
 30484 current set that are not in the set pointed to by the argument *set*; that is, remove each  
 30485 element of the specified set from the current filter.

30486 **RETURN VALUE**

30487 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 30488 return the corresponding error number.

30489 The *posix\_trace\_get\_filter()* function stores the set of filtered trace event types in *set*, if successful.

30490 **ERRORS**

30491 These functions may fail if:

30492 [EINVAL] The value of the *trid* argument does not correspond to an active trace stream  
 30493 or the value of the argument pointed to by *set* is invalid.

30494 [EINTR] The operation was interrupted by a signal.

30495 **EXAMPLES**

30496           None.

30497 **APPLICATION USAGE**

30498           None.

30499 **RATIONALE**

30500           None.

30501 **FUTURE DIRECTIONS**

30502           None.

30503 **SEE ALSO**

30504           *posix\_trace\_eventset\_add()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**trace.h**>

30505 **CHANGE HISTORY**

30506           First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

30507 **NAME**

30508        posix\_trace\_get\_status — retrieve the trace statuses

30509 **SYNOPSIS**

30510 TRC     #include &lt;trace.h&gt;

30511        int posix\_trace\_get\_status(trace\_id\_t trid,  
30512            struct posix\_trace\_status\_info \*statusinfo);

30513

30514 **DESCRIPTION**30515        Refer to *posix\_trace\_get\_attr()*.

30516 NAME

30517 posix\_trace\_getnext\_event, posix\_trace\_timedgetnext\_event, posix\_trace\_trygetnext\_event —  
 30518 retrieve a trace event

30519 SYNOPSIS

```

30520 TRC #include <sys/types.h>
30521 #include <trace.h>

30522 int posix_trace_getnext_event(trace_id_t trid,
30523 struct posix_trace_event_info *event, void *data,
30524 size_t num_bytes, size_t *data_len, int *unavailable);
30525 TRC TMO int posix_trace_timedgetnext_event(trace_id_t trid,
30526 struct posix_trace_event_info *event, void *data,
30527 size_t num_bytes, size_t *data_len, int *unavailable,
30528 const struct timespec *abs_timeout);
30529 int posix_trace_trygetnext_event(trace_id_t trid,
30530 struct posix_trace_event_info *event, void *data,
30531 size_t num_bytes, size_t *data_len, int *unavailable);
30532

```

30533 DESCRIPTION

30534 The *posix\_trace\_getnext\_event()* function shall be used to report a recorded trace event either  
 30535 TRL from an active trace stream without log or a pre-recorded trace stream identified by the *trid*  
 30536 argument. The *posix\_trace\_trygetnext\_event()* function shall be used to report a recorded trace  
 30537 event from an active trace stream without log identified by the *trid* argument.

30538 The trace event information associated with the recorded trace event shall be copied by the  
 30539 function into the structure pointed to by the argument *event* and the data associated with the  
 30540 trace event shall be copied into the buffer pointed to by the *data* argument.

30541 The *posix\_trace\_getnext\_event()* function shall block if the *trid* argument identifies an active trace  
 30542 stream and there is currently no trace event ready to be retrieved. When returning, if a recorded  
 30543 trace event was reported, the variable pointed to by the *unavailable* argument shall be set to zero.  
 30544 Otherwise, the variable pointed to by the *unavailable* argument shall be set to a value different  
 30545 from zero.

30546 TMO The *posix\_trace\_timedgetnext\_event()* function attempts to get another trace event from an active  
 30547 trace stream without log, as in the *posix\_trace\_getnext\_event()* function. However, if no trace  
 30548 event is available from the trace stream, this wait shall be terminated when the timeout specified  
 30549 by the argument *abs\_timeout* expires, and the function shall return the error [ETIMEDOUT].

30550 The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the  
 30551 clock upon which timeouts are based (that is, when the value of that clock equals or exceeds  
 30552 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already passed at the time of the  
 30553 call.

30554 TMO TMR If the Timers option is supported, the timeout is based on the CLOCK\_REALTIME clock; if the  
 30555 Timers option is not supported, the timeout is based on the system clock as returned by the  
 30556 *time()* function. The resolution of the timeout is the resolution of the clock on which it is based.  
 30557 The **timespec** data type is defined as a structure in the header **<time.h>**.

30558 Under no circumstance will the function fail with a timeout if a trace event is immediately  
 30559 available from the trace stream. The validity of the *abs\_timeout* argument need not be checked if  
 30560 a trace event is immediately available from the trace stream.

30561 TMO TMR The behavior of this function for a pre-recorded trace stream is unspecified.

30562 TRL The *posix\_trace\_trygetnext\_event()* function shall not block. This function shall return an error if  
 30563 the *trid* argument identifies a pre-recorded trace stream. If a recorded trace event was reported,  
 30564 the variable pointed to by the *unavailable* argument shall be set to zero. Otherwise, if no trace  
 30565 event was reported, the variable pointed to by the *unavailable* argument shall be set to a value  
 30566 different from zero.

30567 The argument *num\_bytes* shall be the size of the buffer pointed to by the *data* argument. The  
 30568 argument *data\_len* reports to the application the length in bytes of the data record just  
 30569 transferred. If *num\_bytes* is greater than or equal to the size of the data associated with the trace  
 30570 event pointed to by the *event* argument, all the recorded data shall be transferred. In this case, the  
 30571 *truncation-status* member of the trace event structure shall be either  
 30572 POSIX\_TRACE\_NOT\_TRUNCATED, if the trace event data was recorded without truncation  
 30573 while tracing, or POSIX\_TRACE\_TRUNCATED\_RECORD, if the trace event data was truncated  
 30574 when it was recorded. If the *num\_bytes* argument is less than the length of recorded trace event  
 30575 data, the data transferred shall be truncated to a length of *num\_bytes*, the value stored in the  
 30576 variable pointed to by *data\_len* shall be equal to *num\_bytes*, and the *truncation-status* member of  
 30577 the *event* structure argument shall be set to POSIX\_TRACE\_TRUNCATED\_READ (see the  
 30578 *posix\_trace\_event\_info()* function).

30579 The report of a trace event shall be sequential starting from the oldest recorded trace event. Trace  
 30580 events shall be reported in the order in which they were generated, up to an implementation-  
 30581 defined time resolution that causes the ordering of trace events occurring very close to each  
 30582 other to be unknown. Once reported, a trace event cannot be reported again from an active trace  
 30583 stream. Once a trace event is reported from an active trace stream without log, the trace stream  
 30584 shall make the resources associated with that trace event available to record future generated  
 30585 trace events.

#### 30586 RETURN VALUE

30587 Upon successful completion, these functions shall return a value of zero. Otherwise, they shall  
 30588 return the corresponding error number.

30589 If successful, these functions store:

- 30590 • The recorded trace event in the object pointed to by *event*
- 30591 • The trace event information associated with the recorded trace event in the object pointed to  
 30592 by *data*
- 30593 • The length of this trace event information in the object pointed to by *data\_len*
- 30594 • The value of zero in the object pointed to by *unavailable*

#### 30595 ERRORS

30596 These functions may fail if:

30597 [EINVAL] The trace stream identifier argument *trid* is invalid.

30598 The *posix\_trace\_getnext\_event()* and *posix\_trace\_timedgetnext\_event()* functions may fail if:

30599 [EINTR] The operation was interrupted by a signal, and so the call had no effect.

30600 The *posix\_trace\_trygetnext\_event()* function may fail if:

30601 [EINVAL] The trace stream identifier argument *trid* does not correspond to an active  
 30602 trace stream.

30603 TMO The *posix\_trace\_timedgetnext\_event()* function may fail if:

30604 [EINVAL] There is no trace event immediately available from the trace stream, and the  
 30605 *timeout* argument is invalid.

30606 [ETIMEDOUT] No trace event was available from the trace stream before the specified  
30607 timeout *timeout* expired.  
30608

30609 **EXAMPLES**

30610 None.

30611 **APPLICATION USAGE**

30612 None.

30613 **RATIONALE**

30614 None.

30615 **FUTURE DIRECTIONS**

30616 None.

30617 **SEE ALSO**

30618 *posix\_trace\_create()*, **posix\_trace\_event\_info Structure**, *posix\_trace\_open()*, the Base Definitions  
30619 volume of IEEE Std. 1003.1-200x, <sys/types.h>, <trace.h>

30620 **CHANGE HISTORY**

30621 First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

30622 **NAME**

30623        posix\_trace\_open — trace log management

30624 **SYNOPSIS**

30625 TCT TRL #include &lt;trace.h&gt;

30626        int posix\_trace\_open(int *file\_desc*, trace\_id\_t \*trid);

30627

30628 **DESCRIPTION**30629        Refer to *posix\_trace\_close()*.

30630 **NAME**

30631        posix\_trace\_rewind — trace log management

30632 **SYNOPSIS**

30633 TCT TRL #include <trace.h>

30634        int posix\_trace\_rewind(trace\_id\_t *trid*);

30635

30636 **DESCRIPTION**

30637        Refer to *posix\_trace\_close()*.

30638 **NAME**

30639        posix\_trace\_set\_filter — set filter of an initialized trace stream

30640 **SYNOPSIS**

30641 TRC TEF #include &lt;trace.h&gt;

30642        int posix\_trace\_set\_filter(trace\_id\_t *trid*,  
30643            const trace\_event\_set\_t \**set*, int *how*);

30644

30645 **DESCRIPTION**30646        Refer to *posix\_trace\_get\_filter()*.

30647 **NAME**

30648        posix\_trace\_shutdown — trace stream shutdown from a process

30649 **SYNOPSIS**

30650 TRC     #include <sys/types.h>

30651        #include <trace.h>

30652        int posix\_trace\_shutdown(trace\_id\_t *trid*);

30653

30654 **DESCRIPTION**

30655        Refer to *posix\_trace\_create()*.

30656 **NAME**

30657            posix\_trace\_start, posix\_trace\_stop — trace start and stop

30658 **SYNOPSIS**

30659 TRC        #include &lt;trace.h&gt;

30660               int posix\_trace\_start(trace\_id\_t trid);

30661               int posix\_trace\_stop (trace\_id\_t trid);

30662

30663 **DESCRIPTION**30664            The *posix\_trace\_start()* and *posix\_trace\_stop()* functions, respectively, start and stop the trace stream identified by the argument *trid*.30666            The effect of calling the *posix\_trace\_start()* function shall be recorded in the trace stream as the POSIX\_TRACE\_START system trace event and the status of the trace stream shall become POSIX\_TRACE\_RUNNING. If the trace stream is in progress when this function is called, the POSIX\_TRACE\_START system trace event shall not be recorded and the trace stream shall continue to run. If the trace stream is full, the POSIX\_TRACE\_START system trace event shall not be recorded and the status of the trace stream shall not be changed.30672            The effect of calling the *posix\_trace\_stop()* function shall be recorded in the trace stream as the POSIX\_TRACE\_STOP system trace event and the status of the trace stream shall become POSIX\_TRACE\_SUSPENDED. If the trace stream is suspended when this function is called, the POSIX\_TRACE\_STOP system trace event shall not be recorded and the trace stream shall remain suspended. If the trace stream is full, the POSIX\_TRACE\_STOP system trace event shall not be recorded and the status of the trace stream shall not be changed.30678 **RETURN VALUE**

30679            Upon successful completion, these functions shall return a value of zero. Otherwise, they shall return the corresponding error number.

30681 **ERRORS**

30682            These functions may fail if:

30683            [EINVAL]            The value of the argument *trid* does not correspond to an active trace stream and thus no trace stream was started or stopped.

30685            [EINTR]            The operation was interrupted by a signal and thus the trace stream was not necessarily started or stopped.

30687 **EXAMPLES**

30688            None.

30689 **APPLICATION USAGE**

30690            None.

30691 **RATIONALE**

30692            None.

30693 **FUTURE DIRECTIONS**

30694            None.

30695 **SEE ALSO**30696            *posix\_trace\_create()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <trace.h>

30697 **CHANGE HISTORY**

30698 First released in Issue 6. Derived from IEEE Std. 1003.1q-2000.

30699 **NAME**30700 `posix_trace_timedgetnext_event`, — retrieve a trace event30701 **SYNOPSIS**30702 TRC TMO `#include <sys/types.h>`30703 `#include <trace.h>`

```
30704 int posix_trace_timedgetnext_event(trace_id_t trid,
30705 struct posix_trace_event_info *event, void *data,
30706 size_t num_bytes, size_t *data_len, int *unavailable,
30707 const struct timespec *abs_timeout);
30708
```

30709 **DESCRIPTION**30710 Refer to `posix_trace_getnext_event()`.

30711 **NAME**

30712        posix\_trace\_trid\_eventid\_open — manipulate trace event type identifier

30713 **SYNOPSIS**

30714 TRC TEF #include <trace.h>

```
30715 int posix_trace_trid_eventid_open(trace_id_t trid,
30716 const char *event_name, trace_eventid_t *event);
```

30717

30718 **DESCRIPTION**

30719        Refer to *posix\_trace\_eventid\_equal()*.

30720 **NAME**

30721        posix\_trace\_trygetnext\_event — retrieve a trace event

30722 **SYNOPSIS**

30723 TRC TMO #include &lt;sys/types.h&gt;

30724        #include &lt;trace.h&gt;

```
30725 int posix_trace_trygetnext_event(trace_id_t trid,
30726 struct posix_trace_event_info *event, void *data,
30727 size_t num_bytes, size_t *data_len, int *unavailable);
```

30728

30729 **DESCRIPTION**30730        Refer to *posix\_trace\_getnext\_event()*.

30731 **NAME**

30732 posix\_typed\_mem\_get\_info — query typed memory information

30733 **SYNOPSIS**

30734 **TYM** #include <sys/mman.h>

```
30735 int posix_typed_mem_get_info(int fildes,
30736 struct posix_typed_mem_info *info);
```

30737

30738 **DESCRIPTION**

30739 The *posix\_typed\_mem\_get\_info()* function returns, in the *posix\_tmi\_length* field of the  
 30740 **posix\_typed\_mem\_info** structure pointed to by *info*, the maximum length which may be  
 30741 successfully allocated by the typed memory object designated by *fildes*. This maximum length  
 30742 shall take into account the flag `POSIX_TYPED_MEM_ALLOCATE` or  
 30743 `POSIX_TYPED_MEM_ALLOCATE_CONTIG` specified when the typed memory object  
 30744 represented by *fildes* was opened. The maximum length is dynamic; therefore, the value returned  
 30745 is valid only while the current mapping of the corresponding typed memory pool remains  
 30746 unchanged.

30747 If *fildes* represents a typed memory object opened with neither the  
 30748 `POSIX_TYPED_MEM_ALLOCATE` flag nor the `POSIX_TYPED_MEM_ALLOCATE_CONTIG`  
 30749 flag specified, the returned value of *info*->*posix\_tmi\_length* is unspecified.

30750 The *posix\_typed\_mem\_get\_info()* function may return additional implementation-defined  
 30751 information in other fields of the **posix\_typed\_mem\_info** structure pointed to by *info*.

30752 If the memory object specified by *fildes* is not a typed memory object, then the behavior of this  
 30753 function is undefined.

30754 **RETURN VALUE**

30755 Upon successful completion, the *posix\_typed\_mem\_get\_info()* function shall return zero;  
 30756 otherwise, the corresponding error status value shall be returned.

30757 **ERRORS**

30758 The *posix\_typed\_mem\_get\_info()* function shall fail if:

30759 `[EBADF]` The *fildes* argument is not a valid open file descriptor.

30760 `[ENODEV]` The *fildes* argument is not connected to a memory object supported by this  
 30761 function.

30762 This function shall not return an error code of `[EINTR]`.

30763 **EXAMPLES**

30764 None.

30765 **APPLICATION USAGE**

30766 None.

30767 **RATIONALE**

30768 An application that needs to allocate a block of typed memory with length dependent upon the  
 30769 amount of memory currently available must either query the typed memory object to obtain the  
 30770 amount available, or repeatedly invoke *mmap()* attempting to guess an appropriate length.  
 30771 While the latter method is existing practice with *malloc()*, it is awkward and imprecise. The  
 30772 *posix\_typed\_mem\_get\_info()* function allows an application to immediately determine available  
 30773 memory. This is particularly important for typed memory objects that may in some cases be  
 30774 scarce resources. Note that when a typed memory pool is a shared resource, some form of  
 30775 mutual exclusion or synchronization may be required while typed memory is being queried and

30776 allocated to prevent race conditions.

30777 The existing *fstat()* function is not suitable for this purpose. We realize that implementations  
30778 may wish to provide other attributes of typed memory objects (for example, alignment  
30779 requirements, page size, and so on). The *fstat()* function returns a structure which is not  
30780 extensible and, furthermore, contains substantial information that is inappropriate for typed  
30781 memory objects.

30782 **FUTURE DIRECTIONS**

30783 None.

30784 **SEE ALSO**

30785 *fstat()*, *mmap()*, *posix\_typed\_mem\_open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |

30786 <**sys/mman.h**>

30787 **CHANGE HISTORY**

30788 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

30789 **NAME**

30790 posix\_typed\_mem\_open — open a typed memory object

30791 **SYNOPSIS**

30792 TYM #include <sys/mman.h>

30793 int posix\_typed\_mem\_open(const char \*name, int oflag, int tflag);

30794

30795 **DESCRIPTION**

30796 The *posix\_typed\_mem\_open()* function establishes a connection between the typed memory object  
 30797 specified by the string pointed to by *name* and a file descriptor. It creates an open file description  
 30798 that refers to the typed memory object and a file descriptor that refers to that open file  
 30799 description. The file descriptor is used by other functions to refer to that typed memory object. It  
 30800 is unspecified whether the name appears in the file system and is visible to other functions that  
 30801 take path names as arguments. The *name* argument shall conform to the construction rules for a  
 30802 path name. If *name* begins with the slash character, then processes calling  
 30803 *posix\_typed\_mem\_open()* with the same value of *name* shall refer to the same typed memory  
 30804 object. If *name* does not begin with the slash character, the effect is implementation-defined. The  
 30805 interpretation of slash characters other than the leading slash character in *name* is  
 30806 implementation-defined.

30807 Each typed memory object supported in a system is identified by a name which specifies not  
 30808 only its associated typed memory pool, but also the path or port by which it is accessed. That is,  
 30809 the same typed memory pool accessed via several different ports has several different  
 30810 corresponding names. The binding between names and typed memory objects is established in  
 30811 an implementation-defined manner. Unlike shared memory objects, there is ordinarily no way  
 30812 for a program to create a typed memory object.

30813 The value of *tflag* determines how the typed memory object behaves when subsequently  
 30814 mapped by calls to *mmap()*. At most, one of the following flags defined in <sys/mman.h> may  
 30815 be specified:

30816 POSIX\_TYPED\_MEM\_ALLOCATE

30817 Allocate on *mmap()*.

30818 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG

30819 Allocate contiguously on *mmap()*.

30820 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE

30821 Map on *mmap()*, without affecting allocatability.

30822 If *tflag* has the flag POSIX\_TYPED\_MEM\_ALLOCATE specified, any subsequent call to *mmap()*  
 30823 using the returned file descriptor shall result in allocation and mapping of typed memory from  
 30824 the specified typed memory pool. The allocated memory may be a contiguous previously  
 30825 unallocated area of the typed memory pool or several non-contiguous previously unallocated  
 30826 areas (mapped to a contiguous portion of the process address space). If *tflag* has the flag  
 30827 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG specified, any subsequent call to *mmap()* using the  
 30828 returned file descriptor shall result in allocation and mapping of a single contiguous previously  
 30829 unallocated area of the typed memory pool (also mapped to a contiguous portion of the process  
 30830 address space). If *tflag* has none of the flags POSIX\_TYPED\_MEM\_ALLOCATE or  
 30831 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG specified, any subsequent call to *mmap()* using the  
 30832 returned file descriptor shall map an application-chosen area from the specified typed memory  
 30833 pool such that this mapped area becomes unavailable for allocation until unmapped by all  
 30834 processes. If *tflag* has the flag POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE specified, any  
 30835 subsequent call to *mmap()* using the returned file descriptor shall map an application-chosen  
 30836 area from the specified typed memory pool without an effect on the availability of that area for

30837 allocation; that is, mapping such an object leaves each byte of the mapped area unallocated if it  
 30838 was unallocated prior to the mapping or allocated if it was allocated prior to the mapping. The  
 30839 appropriate privilege to specify the POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE flag is  
 30840 implementation-defined.

30841 If successful, *posix\_typed\_mem\_open()* returns a file descriptor for the typed memory object that  
 30842 is the lowest numbered file descriptor not currently open for that process. The open file  
 30843 description is new, and therefore the file descriptor does not share it with any other processes. It  
 30844 is unspecified whether the file offset is set. The FD\_CLOEXEC file descriptor flag associated  
 30845 with the new file descriptor shall be cleared.

30846 The behavior of *msync()*, *ftruncate()*, and all file operations other than *mmap()*,  
 30847 *posix\_mem\_offset()*, *posix\_typed\_mem\_get\_info()*, *fstat()*, *dup()*, *dup2()*, and *close()*, is unspecified  
 30848 when passed a file descriptor connected to a typed memory object by this function.

30849 The file status flags of the open file description shall be set according to the value of *oflag*.  
 30850 Applications shall specify exactly one of the three access mode values described below and  
 30851 defined in the header `<fcntl.h>`, as the value of *oflag*.

30852 O\_RDONLY      Open for read access only.  
 30853 O\_WRONLY      Open for write access only.  
 30854 O\_RDWR        Open for read or write access.

#### 30855 RETURN VALUE

30856 Upon successful completion, the *posix\_typed\_mem\_open()* function shall return a non-negative  
 30857 integer representing the lowest numbered unused file descriptor. Otherwise, it shall return `-1`  
 30858 and set *errno* to indicate the error.

#### 30859 ERRORS

30860 The *posix\_typed\_mem\_open()* function shall fail if:

30861 [EACCES]      The typed memory object exists and the permissions specified by *oflag* are  
 30862 denied.

30863 [EINTR]        The *posix\_typed\_mem\_open()* operation was interrupted by a signal.

30864 [EINVAL]       The flags specified in *tflag* are invalid (more than one of  
 30865 POSIX\_TYPED\_MEM\_ALLOCATE,  
 30866 POSIX\_TYPED\_MEM\_ALLOCATE\_CONTIG, or  
 30867 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE is specified).

30868 [EMFILE]       Too many file descriptors are currently in use by this process.

30869 [ENAMETOOLONG]   The length of the *name* argument exceeds {PATH\_MAX} or a path name  
 30870 component is longer than {NAME\_MAX}.  
 30871

30872 [ENFILE]        Too many file descriptors are currently open in the system.

30873 [ENOENT]        The named typed memory object does not exist.

30874 [EPERM]        The caller lacks the appropriate privilege to specify the flag  
 30875 POSIX\_TYPED\_MEM\_MAP\_ALLOCATABLE in argument *tflag*.

30876 **EXAMPLES**

30877           None.

30878 **APPLICATION USAGE**

30879           None.

30880 **RATIONALE**

30881           None.

30882 **FUTURE DIRECTIONS**

30883           None.

30884 **SEE ALSO**

30885           *close()*, *dup()*, *exec*, *fcntl()*, *fstat()*, *ftruncate()*, *mmap()*, *msync()*, *posix\_mem\_offset()*,  
30886           *posix\_typed\_mem\_get\_info()*, *umask()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
30887           <fcntl.h,> <sys/mman.h>

30888 **CHANGE HISTORY**

30889           First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

30890 **NAME**

30891 pow, powf, powl — power function

30892 **SYNOPSIS**

30893 #include &lt;math.h&gt;

30894 double pow(double x, double y);

30895 float powf(float x, float y);

30896 long double powl(long double x, long double y);

30897 **DESCRIPTION**

30898 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 30899 conflict between the requirements described here and the ISO C standard is unintentional. This  
 30900 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

30901 These functions shall compute the value of  $x$  raised to the power  $y$ ,  $x^y$ . If  $x$  is negative, the  
 30902 application shall ensure that  $y$  is an integer value.

30903 An application wishing to check for error situations should set *errno* to 0 before calling *pow()*. If  
 30904 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

30905 **RETURN VALUE**30906 Upon successful completion, these functions shall return the value of  $x$  raised to the power  $y$ .30907 If  $x$  is 0 and  $y$  is 0, 1.0 shall be returned.

30908 **XSI** If  $y$  is NaN, or  $y$  is non-zero and  $x$  is NaN, NaN shall be returned and *errno* may be set to  
 30909 [EDOM]. If  $y$  is 0.0 and  $x$  is NaN, either 1.0 shall be returned, or NaN shall be returned and *errno*  
 30910 may be set to [EDOM].

30911 **XSI** If  $x$  is 0.0 and  $y$  is negative,  $-\text{HUGE\_VAL}$  shall be returned and *errno* may be set to [EDOM] or  
 30912 [ERANGE].

30913 If the correct value would cause overflow,  $\pm\text{HUGE\_VAL}$  shall be returned, and *errno* set to  
 30914 [ERANGE].

30915 If the correct value would cause underflow, 0 is returned and *errno* may be set to [ERANGE].30916 **ERRORS**

30917 These functions shall fail if:

30918 [EDOM] The value of  $x$  is negative and  $y$  is non-integral.

30919 [ERANGE] The value to be returned would have caused overflow.

30920 These functions may fail if:

30921 **XSI** [EDOM] The value of  $x$  is 0.0 and  $y$  is negative, or  $y$  is NaN.

30922 [ERANGE] The correct value would cause underflow.

30923 **XSI** No other errors shall occur.

30924 **EXAMPLES**

30925 None.

30926 **APPLICATION USAGE**

30927 None.

30928 **RATIONALE**

30929 None.

30930 **FUTURE DIRECTIONS**

30931 None.

30932 **SEE ALSO**30933 *exp()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>30934 **CHANGE HISTORY**

30935 First released in Issue 1. Derived from Issue 1 of the SVID.

30936 **Issue 4**30937 References to *matherr()* are removed.30938 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
30939 ISO C standard and to rationalize error handling in the mathematics functions.

30940 The return value specified for [EDOM] is marked as an extension.

30941 **Issue 5**30942 The DESCRIPTION is updated to indicate how an application should check for an error. This  
30943 text was previously published in the APPLICATION USAGE section.30944 **Issue 6**

30945 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

30946 The *powf()* and *powl()* functions are added for alignment with the ISO/IEC 9899: 1999 standard.

30947 **NAME**

30948        pread — read from a file

30949 **SYNOPSIS**

30950 xSI        #include &lt;unistd.h&gt;

30951        ssize\_t pread(int *fd*, void \**buf*, size\_t *nbyte*, off\_t *offset*);

30952

30953 **DESCRIPTION**30954        Refer to *read()*.

30955 **NAME**

30956       printf — print formatted output

30957 **SYNOPSIS**

30958       #include <stdio.h>

30959       int printf(const char \*restrict *format*, ...);

30960 **DESCRIPTION**

30961       Refer to *fprintf()*.

## 30962 NAME

30963 pselect, select — synchronous I/O multiplexing

## 30964 SYNOPSIS

```

30965 #include <sys/select.h>

30966 int pselect(int nfds, fd_set *readfds, fd_set *writefds,
30967 fd_set *errorfds, const struct timespec *timeout,
30968 const sigset_t *sigmask);
30969 int select(int nfds, fd_set *restrict readfds,
30970 fd_set *restrict writefds, fd_set *restrict errorfds,
30971 struct timeval *restrict timeout);
30972 void FD_CLR(int fd, fd_set *fdset);
30973 int FD_ISSET(int fd, fd_set *fdset);
30974 void FD_SET(int fd, fd_set *fdset);
30975 void FD_ZERO(fd_set *fdset);

```

## 30976 DESCRIPTION

30977 The *pselect()* function shall examine the file descriptor sets whose addresses are passed in the  
 30978 *readfds*, *writefds*, and *errorfds* parameters to see if some of their descriptors are ready for reading,  
 30979 are ready for writing, or have an exceptional condition pending, respectively.

30980 The *select()* function shall be equivalent to the *pselect()* function, except as follows:

- 30981 • For the *select()* function, the timeout period is given in seconds and microseconds in an  
 30982 argument of type **struct timeval**, whereas for the *pselect()* function the timeout period is  
 30983 given in seconds and nanoseconds in an argument of type **struct timespec**.
- 30984 • The *select()* function has no *sigmask* argument; it shall behave as *pselect()* does when *sigmask*  
 30985 is a null pointer.
- 30986 • Upon successful completion, the *select()* function may modify the object pointed to by the  
 30987 *timeout* argument.

30988 The *pselect()* and *select()* functions shall support regular files, terminal and pseudo-terminal  
 30989 XSR devices, **STREAMS**-based files, FIFOs, pipes, and sockets. The behavior of *pselect()* and *select()*  
 30990 on file descriptors that refer to other types of file is unspecified.

30991 The *nfds* argument specifies the range of descriptors to be tested. The first *nfds* descriptors shall  
 30992 be checked in each set; that is, the descriptors from zero through *nfds*-1 in the descriptor sets  
 30993 shall be examined.

30994 If the *readfds* argument is not a null pointer, it points to an object of type **fd\_set** that on input  
 30995 specifies the file descriptors to be checked for being ready to read, and on output indicates  
 30996 which file descriptors are ready to read.

30997 If the *writefds* argument is not a null pointer, it points to an object of type **fd\_set** that on input  
 30998 specifies the file descriptors to be checked for being ready to write, and on output indicates  
 30999 which file descriptors are ready to write.

31000 If the *errorfds* argument is not a null pointer, it points to an object of type **fd\_set** that on input  
 31001 specifies the file descriptors to be checked for error conditions pending, and on output indicates  
 31002 which file descriptors have error conditions pending.

31003 Upon successful completion, the *pselect()* function shall modify the objects pointed to by the  
 31004 *readfds*, *writefds*, and *errorfds* arguments to indicate which file descriptors are ready for reading,  
 31005 ready for writing, or have an error condition pending, respectively, and shall return the total  
 31006 number of ready descriptors in all the output sets. For each file descriptor less than *nfds*, the  
 31007 corresponding bit shall be set on successful completion if it was set on input and the associated

31008 condition is true for that file descriptor.

31009 If none of the selected descriptors are ready for the requested operation, the *pselect()* function  
31010 may block until at least one of the requested operations becomes ready. The parameter *timeout*  
31011 controls how long the *pselect()* function may take to complete. If the *timeout* parameter is not a  
31012 null pointer, it specifies a maximum interval to wait for the selection to complete. If the specified  
31013 time interval expires without any requested operation becoming ready, the function shall return.  
31014 If the *timeout* parameter is a null pointer, then the call to *pselect()* shall block indefinitely until at  
31015 least one descriptor meets the specified criteria. To effect a poll, the *timeout* parameter should not  
31016 be a null pointer, and should point to a zero-valued **timespec** structure.

31017 The use of a timeout does not affect any pending timers set up by *alarm()*, *ualarm()*, or  
31018 *setitimer()*.

31019 Implementations may place limitations on the maximum timeout interval supported. All  
31020 implementations shall support a maximum timeout interval of at least 31 days. If the *timeout*  
31021 argument specifies a timeout interval greater than the implementation-defined maximum value,  
31022 the maximum value shall be used as the actual timeout value. Implementations may also place  
31023 limitations on the granularity of timeout intervals. If the requested timeout interval requires a  
31024 finer granularity than the implementation supports, the actual timeout interval shall be rounded  
31025 up to the next supported value.

31026 If *sigmask* is not a null pointer, then the *pselect()* function shall replace the signal mask of the  
31027 process by the set of signals pointed to by *sigmask* before examining the descriptors, and shall  
31028 restore the signal mask of the process before returning.

31029 A descriptor shall be considered ready for reading when a call to an input function with  
31030 O\_NONBLOCK clear would not block, whether or not the function would transfer data  
31031 successfully. (The function might return data, an end-of-file indication, or an error other than  
31032 one indicating that it is blocked, and in each of these cases the descriptor shall be considered  
31033 ready for reading.)

31034 A descriptor shall be considered ready for writing when a call to an output function with  
31035 O\_NONBLOCK clear would not block, whether or not the function would transfer data  
31036 successfully.

31037 If a socket has a pending error, it shall be considered to have an exceptional condition pending.  
31038 Otherwise, what constitutes an exceptional condition is file type-specific. For a file descriptor for  
31039 use with a socket, it is protocol-specific except as noted below. For other file types it is  
31040 implementation-defined. If the operation is meaningless for a particular file type, *pselect()* shall  
31041 indicate that the descriptor is ready for read or write operations, and shall indicate that the  
31042 descriptor has no exceptional condition pending.

31043 If a descriptor refers to a socket, the implied input function is the *recvmsg()* function with  
31044 parameters requesting normal and ancillary data, such that the presence of either type shall  
31045 cause the socket to be marked as readable. The presence of out-of-band data will be checked if  
31046 the socket option SO\_OOBINLINE has been enabled, as out-of-band data is enqueued with  
31047 normal data. If the socket is currently listening, then it will be marked as readable if an incoming  
31048 connection request has been received, and a call to the *accept()* function is guaranteed to  
31049 complete without blocking.

31050 If a descriptor refers to a socket, the implied output function is the *sendmsg()* function supplying  
31051 an amount of normal data equal to the current value of the SO\_SNDLOWAT option for the  
31052 socket. If a non-blocking call to the *connect()* function has been made for a socket, and the  
31053 connection attempt has either succeeded or failed leaving a pending error, the socket shall be  
31054 marked as writable.

31055 A socket shall be considered to have an exceptional condition pending if a receive operation  
 31056 with `O_NONBLOCK` clear for the open file description and with the `MSG_OOB` flag set would  
 31057 return out-of-band data without blocking. (It is protocol-specific whether the `MSG_OOB` flag  
 31058 would be used to read out-of-band data.) A socket shall also be considered to have an  
 31059 exceptional condition pending if an out-of-band data mark is present in the receive queue. Other  
 31060 circumstances under which a socket may be considered to have an exceptional condition  
 31061 pending are protocol-specific and implementation-defined.

31062 If the `readfds`, `writfds`, and `errorfds` arguments are all null pointers and the `timeout` argument is not  
 31063 a null pointer, `select()` blocks for the time specified, or until interrupted by a signal. If the `readfds`,  
 31064 `writfds`, and `errorfds` arguments are all null pointers and the `timeout` argument is a null pointer,  
 31065 `pselect()` blocks until interrupted by a signal.

31066 File descriptors associated with regular files shall always select true for ready to read, ready to  
 31067 write, and error conditions.

31068 On failure, the objects pointed to by the `readfds`, `writfds`, and `errorfds` arguments are not modified.  
 31069 If the timeout interval expires without the specified condition being true for any of the specified  
 31070 file descriptors, the objects pointed to by the `readfds`, `writfds`, and `errorfds` arguments have all bits  
 31071 set to 0.

31072 File descriptor masks of type `fd_set` can be initialized and tested with `FD_CLR()`, `FD_ISSET()`,  
 31073 `FD_SET()`, and `FD_ZERO()`. It is unspecified whether each of these is a macro or a function. If a  
 31074 macro definition is suppressed in order to access an actual function, or a program defines an  
 31075 external identifier with any of these names, the behavior is undefined.

31076 The macro `FD_CLR(fd, fdsetp)` shall remove the file descriptor `fd` from the set pointed to by `fdsetp`.  
 31077 If `fd` is not a member of this set, there shall be no effect on the set, nor will an error be returned.

31078 The macro `FD_ISSET(fd, fdsetp)` shall evaluate to non-zero if the file descriptor `fd` is a member of  
 31079 the set pointed to by `fdsetp`, and shall evaluate to zero otherwise.

31080 The macro `FD_SET(fd, fdsetp)` shall add the file descriptor `fd` to the set pointed to by `fdsetp`. If the  
 31081 file descriptor `fd` is already in this set, there shall be no effect on the set, nor will an error be  
 31082 returned.

31083 The macro `FD_ZERO(fdsetp)` shall initialize the descriptor set pointed to by `fdsetp` to the null set.  
 31084 No error is returned if the set is not empty at the time `FD_ZERO()` is invoked.

31085 The behavior of these macros is undefined if the `fd` argument is less than 0 or greater than or  
 31086 equal to `FD_SETSIZE`, or if `fd` is not a valid file descriptor, or if any of the arguments are  
 31087 expressions with side effects.

#### 31088 RETURN VALUE

31089 Upon successful completion, the `pselect()` and `select()` functions shall return the total number of  
 31090 bits set in the bit masks. Otherwise, `-1` shall be returned, and shall set `errno` to indicate the error.

31091 `FD_CLR()`, `FD_SET()`, and `FD_ZERO()` return no value. `FD_ISSET()` returns a non-zero value if  
 31092 the bit for the file descriptor `fd` is set in the file descriptor set pointed to by `fdset`, and 0 otherwise.

#### 31093 ERRORS

31094 Under the following conditions, `pselect()` and `select()` shall fail and set `errno` to:

31095 [EBADF] One or more of the file descriptor sets specified a file descriptor that is not a  
 31096 valid open file descriptor.

31097 [EINTR] The function was interrupted before any of the selected events occurred and  
 31098 before the timeout interval expired.

- 31099 XSI If SA\_RESTART has been set for the interrupting signal, it is implementation-  
31100 defined whether the function restarts or returns with [EINTR].
- 31101 [EINVAL] An invalid timeout interval was specified.
- 31102 [EINVAL] The *nfds* argument is less than 0 or greater than FD\_SETSIZE.
- 31103 XSR [EINVAL] One of the specified file descriptors refers to a STREAM or multiplexer that is  
31104 linked (directly or indirectly) downstream from a multiplexer.
- 31105 **EXAMPLES**
- 31106 None.
- 31107 **APPLICATION USAGE**
- 31108 None.
- 31109 **RATIONALE**
- 31110 In previous versions of the Single UNIX Specification, the *select()* function was defined in the  
31111 <sys/time.h> header. This is now changed to <sys/select.h>. The rationale for this change was  
31112 as follows: the introduction of the *pselect()* function included the <sys/select.h> header and the  
31113 <sys/select.h> header defines all the related definitions for the *pselect()* and *select()* functions.  
31114 Backwards-compatibility to existing XSI implementations is handled by allowing <sys/time.h>  
31115 to include <sys/select.h>.
- 31116 **FUTURE DIRECTIONS**
- 31117 None.
- 31118 **SEE ALSO**
- 31119 *accept()*, *alarm()*, *connect()*, *fcntl()*, *poll()*, *read()*, *recvmsg()*, *sendmsg()*, *setitimer()*, *ualarm()*,  
31120 *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/select.h>, <sys/time.h>
- 31121 **CHANGE HISTORY**
- 31122 First released in Issue 4, Version 2.
- 31123 **Issue 5**
- 31124 Moved from X/OPEN UNIX extension to BASE.
- 31125 In the ERRORS section, the text has been changed to indicate that [EINVAL] is returned when  
31126 *nfds* is less than 0 or greater than FD\_SETSIZE. It previously stated less than 0, or greater than or  
31127 equal to FD\_SETSIZE.
- 31128 Text about *timeout* is moved from the APPLICATION USAGE section to the DESCRIPTION.
- 31129 **Issue 6**
- 31130 The Open Group corrigenda item U026/6 has been applied, changing the occurrences of *readfs*  
31131 and *writefs* in the *select()* DESCRIPTION to be *readfds* and *writefds*.
- 31132 Text referring to sockets is added to the DESCRIPTION.
- 31133 The DESCRIPTION and ERRORS sections are updated so that references to STREAMS are  
31134 marked as part of the XSI STREAMS Option Group.
- 31135 The following new requirements on POSIX implementations derive from alignment with the  
31136 Single UNIX Specification:
- 31137 • These functions are now mandatory.
- 31138 The *pselect()* function is added for alignment with IEEE Std. 1003.1g-2000 and additional detail  
31139 related to sockets semantics is added to the DESCRIPTION.
- 31140 The *select()* function now requires inclusion of <sys/select.h>.

31141 The **restrict** keyword is added to the *select()* prototype for alignment with the  
31142 ISO/IEC 9899:1999 standard.

31143 **NAME**

31144 pthread\_atfork — register fork handlers

31145 **SYNOPSIS**

31146 THR #include &lt;pthread.h&gt;

31147 #include &lt;sys/types.h&gt;

31148 #include &lt;unistd.h&gt;

```
31149 int pthread_atfork(void (*prepare)(void), void (*parent)(void),
31150 void (*child)(void));
```

31151

31152 **DESCRIPTION**

31153 The *pthread\_atfork()* function shall declare fork handlers to be called before and after *fork()*, in  
 31154 the context of the thread that called *fork()*. The *prepare* fork handler shall be called before *fork()*  
 31155 processing commences. The *parent* fork handle shall be called after *fork()* processing completes  
 31156 in the parent process. The *child* fork handler shall be called after *fork()* processing completes in  
 31157 the child process. If no handling is desired at one or more of these three points, the  
 31158 corresponding fork handler address(es) may be set to NULL.

31159 The order of calls to *pthread\_atfork()* is significant. The *parent* and *child* fork handlers shall be  
 31160 called in the order in which they were established by calls to *pthread\_atfork()*. The *prepare* fork  
 31161 handlers shall be called in the opposite order.

31162 **RETURN VALUE**

31163 Upon successful completion, *pthread\_atfork()* shall return a value of zero; otherwise, an error  
 31164 number shall be returned to indicate the error.

31165 **ERRORS**31166 The *pthread\_atfork()* function shall fail if:

31167 [ENOMEM] Insufficient table space exists to record the fork handler addresses.

31168 The *pthread\_atfork()* function shall not return an error code of [EINTR].31169 **EXAMPLES**

31170 None.

31171 **APPLICATION USAGE**

31172 None.

31173 **RATIONALE**

31174 There are at least two serious problems with the semantics of *fork()* in a multi-threaded  
 31175 program. One problem has to do with state (for example, memory) covered by mutexes.  
 31176 Consider the case where one thread has a mutex locked and the state covered by that mutex is  
 31177 inconsistent while another thread calls *fork()*. In the child, the mutex is in the locked state  
 31178 (locked by a nonexistent thread and thus can never be unlocked). Having the child simply  
 31179 reinitialize the mutex is unsatisfactory since this approach does not resolve the question about  
 31180 how to correct or otherwise deal with the inconsistent state in the child.

31181 It is suggested that programs that use *fork()* call an *exec* function very soon afterwards in the  
 31182 child process, thus resetting all states. In the meantime, only a short list of async-signal-safe  
 31183 library routines are promised to be available.

31184 Unfortunately, this solution does not address the needs of multi-threaded libraries. Application  
 31185 programs may not be aware that a multi-threaded library is in use, and they feel free to call any  
 31186 number of library routines between the *fork()* and *exec* calls, just as they always have. Indeed,  
 31187 they may be extant single-threaded programs and cannot, therefore, be expected to obey new  
 31188 restrictions imposed by the threads library.

31189 On the other hand, the multi-threaded library needs a way to protect its internal state during  
 31190 *fork()* in case it is re-entered later in the child process. The problem arises especially in multi-  
 31191 threaded I/O libraries, which are almost sure to be invoked between the *fork()* and *exec* calls to  
 31192 effect I/O redirection. The solution may require locking mutex variables during *fork()*, or it may  
 31193 entail simply resetting the state in the child after the *fork()* processing completes.

31194 The *pthread\_atfork()* function provides multi-threaded libraries with a means to protect  
 31195 themselves from innocent application programs that call *fork()*, and it provides multi-threaded  
 31196 application programs with a standard mechanism for protecting themselves from *fork()* calls in  
 31197 a library routine or the application itself.

31198 The expected usage is that the *prepare* handler acquires all mutex locks and the other two fork  
 31199 handlers release them.

31200 For example, an application can supply a *prepare* routine that acquires the necessary mutexes the  
 31201 library maintains and supply *child* and *parent* routines that release those mutexes, thus ensuring  
 31202 that the child gets a consistent snapshot of the state of the library (and that no mutexes are left  
 31203 stranded). Alternatively, some libraries might be able to supply just a *child* routine that re-  
 31204 initializes the mutexes in the library and all associated states to some known value (for example,  
 31205 what it was when the image was originally executed).

31206 When *fork()* is called, only the calling thread is duplicated in the child process. Synchronization  
 31207 variables remain in the same state in the child as they were in the parent at the time *fork()* was  
 31208 called. Thus, for example, mutex locks may be held by threads that no longer exist in the child  
 31209 process, and any associated states may be inconsistent. The parent process may avoid this by  
 31210 explicit code that acquires and releases locks critical to the child via *pthread\_atfork()*. In addition,  
 31211 any critical threads need to be recreated and re-initialized to the proper state in the child (also  
 31212 via *pthread\_atfork()*).

31213 A higher-level package may acquire locks on its own data structures before invoking lower-level  
 31214 packages. Under this scenario, the order specified for fork handler calls allows a simple rule of  
 31215 initialization for avoiding package deadlock: a package initializes all packages on which it  
 31216 depends before it calls the *pthread\_atfork()* function for itself.

#### 31217 **FUTURE DIRECTIONS**

31218 None.

#### 31219 **SEE ALSO**

31220 *atexit()*, *fork()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**sys/types.h**>

#### 31221 **CHANGE HISTORY**

31222 First released in Issue 5. Derived from the POSIX Threads Extension.

31223 IEEE PASC Interpretation 1003.1c #4 is included.

#### 31224 **Issue 6**

31225 The *pthread\_atfork()* function is marked as part of the Threads option.

31226 The <**pthread.h**> header is added to the SYNOPSIS.

31227 **NAME**

31228 pthread\_attr\_destroy, pthread\_attr\_init — destroy and initialize threads attributes object

31229 **SYNOPSIS**

```
31230 THR #include <pthread.h>
```

```
31231 int pthread_attr_destroy(pthread_attr_t *attr);
```

```
31232 int pthread_attr_init(pthread_attr_t *attr);
```

31233

31234 **DESCRIPTION**

31235 The *pthread\_attr\_destroy()* function is used to destroy a thread attributes object. An  
31236 implementation may cause *pthread\_attr\_destroy()* to set *attr* to an implementation-defined  
31237 invalid value. The behavior of using the attribute after it has been destroyed is undefined.

31238 The *pthread\_attr\_init()* function initializes a thread attributes object *attr* with the default value  
31239 for all of the individual attributes used by a given implementation.

31240 The resulting attributes object (possibly modified by setting individual attribute values), when  
31241 used by *pthread\_create()* defines the attributes of the thread created. A single attributes object can  
31242 be used in multiple simultaneous calls to *pthread\_create()*. The effect of initializing an already  
31243 initialized thread attributes object is undefined.

31244 **RETURN VALUE**

31245 Upon successful completion, *pthread\_attr\_destroy()* and *pthread\_attr\_init()* shall return a value of  
31246 0; otherwise, an error number shall be returned to indicate the error.

31247 **ERRORS**

31248 The *pthread\_attr\_init()* function shall fail if:

31249 [ENOMEM] Insufficient memory exists to initialize the thread attributes object.

31250 These functions shall not return an error code of [EINTR].

31251 **EXAMPLES**

31252 None.

31253 **APPLICATION USAGE**

31254 None.

31255 **RATIONALE**

31256 Attributes objects are provided for threads, mutexes, and condition variables as a mechanism to  
31257 support probable future standardization in these areas without requiring that the function itself  
31258 be changed.

31259 Attributes objects provide clean isolation of the configurable aspects of threads. For example,  
31260 “stack size” is an important attribute of a thread, but it cannot be expressed portably. When  
31261 porting a threaded program, stack sizes often need to be adjusted. The use of attributes objects  
31262 can help by allowing the changes to be isolated in a single place, rather than being spread across  
31263 every instance of thread creation.

31264 Attributes objects can be used to set up “classes” of threads with similar attributes; for example,  
31265 “threads with large stacks and high priority” or “threads with minimal stacks”. These classes  
31266 can be defined in a single place and then referenced wherever threads need to be created.  
31267 Changes to “class” decisions become straightforward, and detailed analysis of each  
31268 *pthread\_create()* call is not required.

31269 The attributes objects are defined as opaque types as an aid to extensibility. If these objects had  
31270 been specified as structures, adding new attributes would force recompilation of all multi-  
31271 threaded programs when the attributes objects are extended; this might not be possible if

31272 different program components were supplied by different vendors.

31273 Additionally, opaque attributes objects present opportunities for improving performance.  
 31274 Argument validity can be checked once when attributes are set, rather than each time a thread is  
 31275 created. Implementations often need to cache kernel objects that are expensive to create.  
 31276 Opaque attributes objects provide an efficient mechanism to detect when cached objects become  
 31277 invalid due to attribute changes.

31278 Because assignment is not necessarily defined on a given opaque type, implementation-defined  
 31279 default values cannot be defined in a portable way. The solution to this problem is to allow  
 31280 attributes objects to be initialized dynamically by attributes object initialization functions, so  
 31281 that default values can be supplied automatically by the implementation.

31282 The following proposal was provided as a suggested alternative to the supplied attributes:

- 31283 1. Maintain the style of passing a parameter formed by the bitwise-inclusive OR of flags to  
 31284 the initialization routines (*pthread\_create()*, *pthread\_mutex\_init()*, *pthread\_cond\_init()*). The  
 31285 parameter containing the flags should be an opaque type for extensibility. If no flags are  
 31286 set in the parameter, then the objects are created with default characteristics. An  
 31287 implementation may specify implementation-defined flag values and associated behavior.
- 31288 2. If further specialization of mutexes and condition variables is necessary, implementations  
 31289 may specify additional procedures that operate on the **pthread\_mutex\_t** and  
 31290 **pthread\_cond\_t** objects (instead of on attributes objects).

31291 The difficulties with this solution are:

- 31292 1. A bitmask is not opaque if bits have to be set into bitvector attributes objects using  
 31293 explicitly-coded bitwise-inclusive OR operations. If the set of options exceeds an **int**,  
 31294 application programmers need to know the location of each bit. If bits are set or read by  
 31295 encapsulation (that is, get or set functions), then the bitmask is merely an implementation  
 31296 of attributes objects as currently defined and should not be exposed to the programmer.
- 31297 2. Many attributes are not Boolean or very small integral values. For example, scheduling  
 31298 policy may be placed in 3-bit or 4-bit, but priority requires 5-bit or more, thereby taking up  
 31299 at least 8-bit out of a possible 16-bit on machines with 16-bit integers. Because of this, the  
 31300 bitmask can only reasonably control whether particular attributes are set or not, and it  
 31301 cannot serve as the repository of the value itself. The value needs to be specified as a  
 31302 function parameter (which is non-extensible), or by setting a structure field (which is non-  
 31303 opaque), or by get and set functions (making the bitmask a redundant addition to the  
 31304 attributes objects).

31305 Stack size is defined as an optional attribute because the very notion of a stack is inherently  
 31306 machine-dependent. Some implementations may not be able to change the size of the stack, for  
 31307 example, and others may not need to because stack pages may be discontinuous and can be  
 31308 allocated and released on demand.

31309 The attribute mechanism has been designed in large measure for extensibility. Future extensions  
 31310 to the attribute mechanism or to any attributes object defined in this volume of  
 31311 IEEE Std. 1003.1-200x has to be done with care so as not to affect binary-compatibility.

31312 Attributes objects, even if allocated by means of dynamic allocation functions such as *malloc()*,  
 31313 may have their size fixed at compile time. This means, for example, a *pthread\_create()* in an  
 31314 implementation with extensions to the **pthread\_attr\_t** cannot look beyond the area that the  
 31315 binary application assumes is valid. This suggests that implementations should maintain a size  
 31316 field in the attributes object, as well as possibly version information, if extensions in different  
 31317 directions (possibly by different vendors) are to be accommodated.

31318 **FUTURE DIRECTIONS**

31319       None.

31320 **SEE ALSO**

31321       *pthread\_attr\_getstackaddr()*,       *pthread\_attr\_getstacksize()*,       *pthread\_attr\_getdetachstate()*,  
31322       *pthread\_create()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**pthread.h**>

31323 **CHANGE HISTORY**

31324       First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31325 **Issue 6**

31326       The *pthread\_attr\_destroy()* and *pthread\_attr\_init()* functions marked as part of the Threads  
31327       option.

31328       IEEE PASC Interpretation 1003.1 #107 is applied, noting that the effect of initializing an already  
31329       initialized thread attributes object is undefined.

31330 **NAME**

31331 pthread\_attr\_getdetachstate, pthread\_attr\_setdetachstate — get or set detachstate attribute

31332 **SYNOPSIS**

```
31333 THR #include <pthread.h>
31334
31334 int pthread_attr_getdetachstate(const pthread_attr_t *attr,
31335 int *detachstate);
31336 int pthread_attr_setdetachstate(pthread_attr_t *attr, int detachstate);
31337
```

31338 **DESCRIPTION**

31339 The *detachstate* attribute controls whether the thread is created in a detached state. If the thread  
 31340 is created detached, then use of the ID of the newly created thread by the *pthread\_detach()* or  
 31341 *pthread\_join()* function is an error.

31342 The *pthread\_attr\_getdetachstate()* and *pthread\_attr\_setdetachstate()* functions, respectively, get and  
 31343 set the *detachstate* attribute in the *attr* object.

31344 For *pthread\_attr\_getdetachstate()*, *detachstate* shall be set to either  
 31345 PTHREAD\_CREATE\_DETACHED or PTHREAD\_CREATE\_JOINABLE.

31346 For *pthread\_attr\_setdetachstate()*, the application shall set *detachstate* to either  
 31347 PTHREAD\_CREATE\_DETACHED or PTHREAD\_CREATE\_JOINABLE.

31348 A value of PTHREAD\_CREATE\_DETACHED causes all threads created with *attr* to be in the  
 31349 detached state, whereas using a value of PTHREAD\_CREATE\_JOINABLE shall cause all threads  
 31350 created with *attr* to be in the joinable state. The default value of the *detachstate* attribute is  
 31351 PTHREAD\_CREATE\_JOINABLE.

31352 **RETURN VALUE**

31353 Upon successful completion, *pthread\_attr\_getdetachstate()* and *pthread\_attr\_setdetachstate()* shall  
 31354 return a value of 0; otherwise, an error number shall be returned to indicate the error.

31355 The *pthread\_attr\_getdetachstate()* function stores the value of the *detachstate* attribute in *detachstate*  
 31356 if successful.

31357 **ERRORS**

31358 The *pthread\_attr\_setdetachstate()* function shall fail if:

31359 [EINVAL] The value of *detachstate* was not valid

31360 These functions shall not return an error code of [EINTR].

31361 **EXAMPLES**

31362 None.

31363 **APPLICATION USAGE**

31364 None.

31365 **RATIONALE**

31366 None.

31367 **FUTURE DIRECTIONS**

31368 None.

31369 **SEE ALSO**

31370 *pthread\_attr\_destroy()*, *pthread\_attr\_getstackaddr()*, *pthread\_attr\_getstacksize()*, *pthread\_create()*, the  
 31371 Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

31372 **CHANGE HISTORY**

31373 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31374 **Issue 6**

31375 The *pthread\_attr\_setdetachstate()* and *pthread\_attr\_getdetachstate()* functions are marked as part of  
31376 the Threads option.

31377 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

31378 **NAME**

31379 pthread\_attr\_getguardsize, pthread\_attr\_setguardsize — get or set the thread guardsize  
 31380 attribute

31381 **SYNOPSIS**

```
31382 xSI #include <pthread.h>
31383
31383 int pthread_attr_getguardsize(const pthread_attr_t *restrict attr,
31384 size_t *restrict guardsize);
31385 int pthread_attr_setguardsize(pthread_attr_t *attr,
31386 size_t guardsize);
31387
```

31388 **DESCRIPTION**

31389 The *guardsize* attribute controls the size of the guard area for the created thread's stack. The  
 31390 *guardsize* attribute provides protection against overflow of the stack pointer. If a thread's stack is  
 31391 created with guard protection, the implementation allocates extra memory at the overflow end  
 31392 of the stack as a buffer against stack overflow of the stack pointer. If an application overflows  
 31393 into this buffer an error results (possibly in a SIGSEGV signal being delivered to the thread).

31394 The *pthread\_attr\_getguardsize()* function gets the *guardsize* attribute in the *attr* object. This  
 31395 attribute is returned in the *guardsize* parameter.

31396 The *pthread\_attr\_setguardsize()* function sets the *guardsize* attribute in the *attr* object. The new  
 31397 value of this attribute is obtained from the *guardsize* parameter. If *guardsize* is zero, a guard area  
 31398 shall not be provided for threads created with *attr*. If *guardsize* is greater than zero, a guard area  
 31399 of at least size *guardsize* bytes is provided for each thread created with *attr*.

31400 A conforming implementation is permitted to round up the value contained in *guardsize* to a  
 31401 multiple of the configurable system variable {PAGESIZE} (see <sys/mman.h>). If an  
 31402 implementation rounds up the value of *guardsize* to a multiple of {PAGESIZE}, a call to  
 31403 *pthread\_attr\_getguardsize()* specifying *attr* shall store in the *guardsize* parameter the guard size  
 31404 specified by the previous *pthread\_attr\_setguardsize()* function call.

31405 The default value of the *guardsize* attribute is {PAGESIZE} bytes. The actual value of {PAGESIZE}  
 31406 is implementation-defined and may not be the same on all implementations.

31407 If the *stackaddr* attribute has been set (that is, the caller is allocating and managing its own thread  
 31408 stacks), the *guardsize* attribute is ignored and no protection shall be provided by the  
 31409 implementation. It is the responsibility of the application to manage stack overflow along with  
 31410 stack allocation and management in this case.

31411 **RETURN VALUE**

31412 If successful, the *pthread\_attr\_getguardsize()* and *pthread\_attr\_setguardsize()* functions shall return  
 31413 zero; otherwise, an error number shall be returned to indicate the error.

31414 **ERRORS**

31415 The *pthread\_attr\_getguardsize()* and *pthread\_attr\_setguardsize()* functions shall fail if:

31416 [EINVAL] The attribute *attr* is invalid.

31417 [EINVAL] The parameter *guardsize* is invalid.

31418 These functions shall not return an error code of [EINTR].

31419 **EXAMPLES**

31420 None.

31421 **APPLICATION USAGE**

31422 None.

31423 **RATIONALE**31424 The *guardsize* attribute is provided to the application for two reasons:

- 31425 1. Overflow protection can potentially result in wasted system resources. An application  
31426 that creates a large number of threads, and which knows its threads never overflow their  
31427 stack, can save system resources by turning off guard areas.
- 31428 2. When threads allocate large data structures on the stack, large guard areas may be needed  
31429 to detect stack overflow.

31430 **FUTURE DIRECTIONS**

31431 None.

31432 **SEE ALSO**31433 The Base Definitions volume of IEEE Std. 1003.1-200x, `<pthread.h>`, `<sys/mman.h>`31434 **CHANGE HISTORY**

31435 First released in Issue 5.

31436 **Issue 6**31437 In the **ERRORS** section, a third [EINVAL] error condition is removed as it is covered by the  
31438 second error condition.31439 The **restrict** keyword is added to the *pthread\_attr\_getguardsize()* prototype for alignment with the  
31440 ISO/IEC 9899:1999 standard.

31441 **NAME**

31442 pthread\_attr\_getinheritsched, pthread\_attr\_setinheritsched — get and set inheritsched attribute  
 31443 (**REALTIME THREADS**)

31444 **SYNOPSIS**

31445 TPS #include <pthread.h>

```
31446 int pthread_attr_getinheritsched(const pthread_attr_t *restrict attr,
31447 int *restrict inheritsched);
31448 int pthread_attr_setinheritsched(pthread_attr_t *attr,
31449 int inheritsched);
31450
```

31451 **DESCRIPTION**

31452 The *pthread\_attr\_getinheritsched()*, and *pthread\_attr\_setinheritsched()* functions, respectively, get  
 31453 and set the *inheritsched* attribute in the *attr* argument.

31454 When the attributes objects are used by *pthread\_create()*, the *inheritsched* attribute determines  
 31455 how the other scheduling attributes of the created thread shall be set:

31456 **PTHREAD\_INHERIT\_SCHED**

31457 Specifies that the scheduling policy and associated attributes shall be inherited from the  
 31458 creating thread, and the scheduling attributes in this *attr* argument shall be ignored.

31459 **PTHREAD\_EXPLICIT\_SCHED**

31460 Specifies that the scheduling policy and associated attributes shall be set to the  
 31461 corresponding values from this attributes object.

31462 The symbols **PTHREAD\_INHERIT\_SCHED** and **PTHREAD\_EXPLICIT\_SCHED** are defined in  
 31463 the header <**pthread.h**>.

31464 **RETURN VALUE**

31465 If successful, the *pthread\_attr\_getinheritsched()* and *pthread\_attr\_setinheritsched()* functions shall  
 31466 return zero; otherwise, an error number shall be returned to indicate the error.

31467 **ERRORS**

31468 The *pthread\_attr\_setinheritsched()* function may fail if:

31469 [EINVAL] The value of *inheritsched* is not valid.

31470 [ENOTSUP] An attempt was made to set the attribute to an unsupported value.

31471 These functions shall not return an error code of [EINTR].

31472 **EXAMPLES**

31473 None.

31474 **APPLICATION USAGE**

31475 After these attributes have been set, a thread can be created with the specified attributes using  
 31476 *pthread\_create()*. Using these routines does not affect the current running thread.

31477 **RATIONALE**

31478 None.

31479 **FUTURE DIRECTIONS**

31480 None.

31481 **SEE ALSO**

31482 *pthread\_attr\_destroy()*, *pthread\_attr\_getscope()*, *pthread\_attr\_getschedpolicy()*,  
31483 *pthread\_attr\_getschedparam()*, *pthread\_create()*, the Base Definitions volume of  
31484 IEEE Std. 1003.1-200x, <pthread.h>, <sched.h>

31485 **CHANGE HISTORY**

31486 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31487 Marked as part of the Realtime Threads Feature Group.

31488 **Issue 6**

31489 The *pthread\_attr\_getinheritsched()* and *pthread\_attr\_setinheritsched()* functions are marked as part  
31490 of the Thread Execution Scheduling option.

31491 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
31492 implementation does not support the Thread Execution Scheduling option.

31493 The **restrict** keyword is added to the *pthread\_attr\_getinheritsched()* prototype for alignment with  
31494 the ISO/IEC 9899:1999 standard.

31495 **NAME**

31496 pthread\_attr\_getschedparam, pthread\_attr\_setschedparam — get and set schedparam attribute

31497 **SYNOPSIS**

31498 THR #include <pthread.h>

```
31499 int pthread_attr_getschedparam(const pthread_attr_t *restrict attr,
31500 struct sched_param *restrict param);
```

```
31501 int pthread_attr_setschedparam(pthread_attr_t *restrict attr,
31502 const struct sched_param *restrict param);
```

31503

31504 **DESCRIPTION**

31505 The *pthread\_attr\_getschedparam()*, and *pthread\_attr\_setschedparam()* functions, respectively, get  
31506 and set the scheduling parameter attributes in the *attr* argument. The contents of the *param*  
31507 structure are defined in <sched.h>. For the SCHED\_FIFO and SCHED\_RR policies, the only  
31508 required member of *param* is *sched\_priority*.

31509 TSP For the SCHED\_SPORADIC policy, the required members of the *param* structure are  
31510 *sched\_priority*, *sched\_ss\_low\_priority*, *sched\_ss\_repl\_period*, *sched\_ss\_init\_budget*, and  
31511 *sched\_ss\_max\_repl*. The specified *sched\_ss\_repl\_period* must be greater than or equal to the  
31512 specified *sched\_ss\_init\_budget* for the function to succeed; if it is not, then the function shall fail.  
31513 The value of *sched\_ss\_max\_repl* shall be within the inclusive range [1,{SS\_REPL\_MAX}] for the  
31514 function to succeed; if not, the function shall fail.

31515 **RETURN VALUE**

31516 If successful, the *pthread\_attr\_getschedparam()* and *pthread\_attr\_setschedparam()* functions shall  
31517 return zero; otherwise, an error number shall be returned to indicate the error.

31518 **ERRORS**

31519 The *pthread\_attr\_setschedparam()* function may fail if:

31520 [EINVAL] The value of *param* is not valid.

31521 [ENOTSUP] An attempt was made to set the attribute to an unsupported value.

31522 These functions shall not return an error code of [EINTR].

31523 **EXAMPLES**

31524 None.

31525 **APPLICATION USAGE**

31526 After these attributes have been set, a thread can be created with the specified attributes using  
31527 *pthread\_create()*. Using these routines does not affect the current running thread.

31528 **RATIONALE**

31529 None.

31530 **FUTURE DIRECTIONS**

31531 None.

31532 **SEE ALSO**

31533 *pthread\_attr\_destroy()*, *pthread\_attr\_getscope()*, *pthread\_attr\_getinheritsched()*,  
31534 *pthread\_attr\_getschedpolicy()*, *pthread\_create()*, the Base Definitions volume of  
31535 IEEE Std. 1003.1-200x, <pthread.h>, <sched.h>

## 31536 CHANGE HISTORY

31537 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

## 31538 Issue 6

31539 The *pthread\_attr\_getschedparam()* and *pthread\_attr\_setschedparam()* functions are marked as part  
31540 of the Threads option.

31541 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std. 1003.1d-1999.

31542 The **restrict** keyword is added to the *pthread\_attr\_getschedparam()* and  
31543 *pthread\_attr\_setschedparam()* prototypes for alignment with the ISO/IEC 9899:1999 standard.

31544 **NAME**

31545 pthread\_attr\_getschedpolicy, pthread\_attr\_setschedpolicy — get and set schedpolicy attribute  
 31546 (**REALTIME THREADS**)

31547 **SYNOPSIS**

```
31548 TPS #include <pthread.h>
31549
31549 int pthread_attr_getschedpolicy(const pthread_attr_t *restrict attr,
31550 int *restrict policy);
31551 int pthread_attr_setschedpolicy(pthread_attr_t *attr, int policy);
31552
```

31553 **DESCRIPTION**

31554 The *pthread\_attr\_getschedpolicy()* and *pthread\_attr\_setschedpolicy()* functions, respectively, get and  
 31555 set the *schedpolicy* attribute in the *attr* argument.

31556 The supported values of *policy* shall include SCHED\_FIFO, SCHED\_RR, and SCHED\_OTHER,  
 31557 which are defined by the header **<sched.h>**. When threads executing with the scheduling policy  
 31558 TSP SCHED\_FIFO, SCHED\_RR, or SCHED\_SPORADIC are waiting on a mutex, they acquire the  
 31559 mutex in priority order when the mutex is unlocked.

31560 **RETURN VALUE**

31561 If successful, the *pthread\_attr\_getschedpolicy()* and *pthread\_attr\_setschedpolicy()* functions shall  
 31562 return zero; otherwise, an error number shall be returned to indicate the error.

31563 **ERRORS**

31564 The *pthread\_attr\_setschedpolicy()* function may fail if:

31565 [EINVAL] The value of *policy* is not valid.  
 31566 [ENOTSUP] An attempt was made to set the attribute to an unsupported value.  
 31567 These functions shall not return an error code of [EINTR].

31568 **EXAMPLES**

31569 None.

31570 **APPLICATION USAGE**

31571 After these attributes have been set, a thread can be created with the specified attributes using  
 31572 *pthread\_create()*. Using these routines does not affect the current running thread.

31573 **RATIONALE**

31574 None.

31575 **FUTURE DIRECTIONS**

31576 None.

31577 **SEE ALSO**

31578 *pthread\_attr\_destroy()*, *pthread\_attr\_getscope()*, *pthread\_attr\_getinheritsched()*,  
 31579 *pthread\_attr\_getschedparam()*, *pthread\_create()*, the Base Definitions volume of  
 31580 IEEE Std. 1003.1-200x, **<pthread.h>**, **<sched.h>**

31581 **CHANGE HISTORY**

31582 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31583 Marked as part of the Realtime Threads Feature Group.

31584 **Issue 6**

31585 The *pthread\_attr\_getschedpolicy()* and *pthread\_attr\_setschedpolicy()* functions are marked as part of  
31586 the Thread Execution Scheduling option.

31587 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
31588 implementation does not support the Thread Execution Scheduling option.

31589 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std. 1003.1d-1999.

31590 The **restrict** keyword is added to the *pthread\_attr\_getschedpolicy()* prototype for alignment with  
31591 the ISO/IEC 9899:1999 standard.

31592 **NAME**

31593 pthread\_attr\_getscope, pthread\_attr\_setscope — get and set contention scope attribute  
 31594 (**REALTIME THREADS**)

31595 **SYNOPSIS**

```
31596 TPS #include <pthread.h>
31597
31597 int pthread_attr_getscope(const pthread_attr_t *restrict attr,
31598 int *restrict contention_scope);
31599 int pthread_attr_setscope(pthread_attr_t *attr, int contention_scope);
31600
```

31601 **DESCRIPTION**

31602 The *pthread\_attr\_getscope()* and *pthread\_attr\_setscope()* functions, respectively, are used to get  
 31603 and set the *contention\_scope* attribute in the *attr* object.

31604 The *contention\_scope* attribute may have the values `PTHREAD_SCOPE_SYSTEM`, signifying  
 31605 system scheduling contention scope, or `PTHREAD_SCOPE_PROCESS`, signifying process  
 31606 scheduling contention scope. The symbols `PTHREAD_SCOPE_SYSTEM` and  
 31607 `PTHREAD_SCOPE_PROCESS` are defined by the header `<pthread.h>`.

31608 **RETURN VALUE**

31609 If successful, the *pthread\_attr\_getscope()* and *pthread\_attr\_setscope()* functions shall return zero;  
 31610 otherwise, an error number shall be returned to indicate the error.

31611 **ERRORS**

31612 The *pthread\_attr\_setscope()* function may fail if:

- 31613 [EINVAL] The value of *contention\_scope* is not valid.
- 31614 [ENOTSUP] An attempt was made to set the attribute to an unsupported value.
- 31615 These functions shall not return an error code of [EINTR].

31616 **EXAMPLES**

31617 None.

31618 **APPLICATION USAGE**

31619 After these attributes have been set, a thread can be created with the specified attributes using  
 31620 *pthread\_create()*. Using these routines does not affect the current running thread.

31621 **RATIONALE**

31622 None.

31623 **FUTURE DIRECTIONS**

31624 None.

31625 **SEE ALSO**

31626 *pthread\_attr\_destroy()*, *pthread\_attr\_getinheritsched()*, *pthread\_attr\_getschedpolicy()*,  
 31627 *pthread\_attr\_getschedparam()*, *pthread\_create()*, the Base Definitions volume of  
 31628 IEEE Std. 1003.1-200x, `<pthread.h>`, `<sched.h>`

31629 **CHANGE HISTORY**

- 31630 First released in Issue 5. Included for alignment with the POSIX Threads Extension.
- 31631 Marked as part of the Realtime Threads Feature Group.

31632 **Issue 6**

31633 The *pthread\_attr\_getscope()* and *pthread\_attr\_setscope()* functions are marked as part of the  
31634 Thread Execution Scheduling option.

31635 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
31636 implementation does not support the Thread Execution Scheduling option.

31637 The **restrict** keyword is added to the *pthread\_attr\_getscope()* prototype for alignment with the  
31638 ISO/IEC 9899:1999 standard.

31639 **NAME**

31640 pthread\_attr\_getstack, pthread\_attr\_setstack — get and set stack attribute

31641 **SYNOPSIS**31642 XSI 

```
#include <pthread.h>
```

31643 

```
int pthread_attr_getstack(const pthread_attr_t *attr, void **stackaddr,
31644 size_t *stacksize);
```

31645 

```
int pthread_attr_setstack(pthread_attr_t *attr, void *stackaddr,
31646 size_t stacksize);
```

31647

31648 **DESCRIPTION**31649 The *pthread\_attr\_getstack()* and *pthread\_attr\_setstack()* functions, respectively, shall get and set  
31650 the thread creation stack attribute in the *attr* object.31651 The *stack* attribute specifies the area of storage to be used for the created thread's stack. The base  
31652 (lowest addressable byte) of the storage is *stackaddr*, and the size of the storage is *stacksize* bytes.  
31653 The *stacksize* shall be at least {PTHREAD\_STACK\_MIN}. The *stackaddr* shall be aligned  
31654 appropriately to be used as a stack; for example, *pthread\_attr\_setstack()* may fail with [EINVAL]  
31655 if (*stackaddr* & 0x7) is not 0. All pages within the stack described by *stackaddr* and *stacksize* shall be  
31656 both readable and writable by the thread.31657 **RETURN VALUE**31658 Upon successful completion, these functions shall return a value of 0; otherwise, an error  
31659 number shall be returned to indicate the error.31660 The *pthread\_attr\_getstack()* function shall store the stack attribute values in *stackaddr* and *stacksize*  
31661 if successful.31662 **ERRORS**31663 The *pthread\_attr\_setstack()* function shall fail if:31664 [EINVAL] The value of *stacksize* is less than {PTHREAD\_STACK\_MIN} or exceeds an  
31665 implementation-defined limit.31666 The *pthread\_attr\_setstack()* function may fail if:31667 [EINVAL] The value of *stackaddr* does not have proper alignment to be used as a stack, or  
31668 if (*stackaddr* + *stacksize*) lacks proper alignment.31669 [EACCES] The stack page(s) described by *stackaddr* and *stacksize* are not both readable  
31670 and writable by the thread.

31671 These functions shall not return an error code of [EINTR].

31672 **EXAMPLES**

31673 None.

31674 **APPLICATION USAGE**31675 These functions are appropriate for use by applications in an environment where the stack for a  
31676 thread must be placed in some particular region of memory.31677 While it might seem that an application could detect stack overflow by providing a protected  
31678 page outside the specified stack region, this cannot be done portably. Implementations are free  
31679 to place the thread's initial stack pointer anywhere within the specified region to accommodate  
31680 the machine's stack pointer behavior and allocation requirements. Furthermore, on some  
31681 architectures, such as the IA-64, "overflow" might mean that two separate stack pointers  
31682 allocated within the region will overlap somewhere in the middle of the region.

31683 **RATIONALE**

31684           None.

31685 **FUTURE DIRECTIONS**

31686           None.

31687 **SEE ALSO**

31688           *pthread\_attr\_init()*, *pthread\_attr\_setdetachstate()*, *pthread\_attr\_setstacksize()*, *pthread\_create()*, the  
31689           Base Definitions volume of IEEE Std. 1003.1-200x, <limits.h>, <pthread.h>

31690 **CHANGE HISTORY**

31691           First released in Issue 6.

31692 **NAME**

31693 pthread\_attr\_getstackaddr, pthread\_attr\_setstackaddr — get and set stackaddr attribute

31694 **SYNOPSIS**

31695 THR TSA #include <pthread.h>

```
31696 int pthread_attr_getstackaddr(const pthread_attr_t *restrict attr,
31697 void **restrict stackaddr);
31698 int pthread_attr_setstackaddr(pthread_attr_t *attr, void *stackaddr);
31699
```

31700 **DESCRIPTION**

31701 The *pthread\_attr\_getstackaddr()* and *pthread\_attr\_setstackaddr()* functions, respectively, get and set  
31702 the thread creation *stackaddr* attribute in the *attr* object.

31703 The *stackaddr* attribute specifies the location of storage to be used for the created thread's stack.  
31704 The size of the storage is at least {PTHREAD\_STACK\_MIN}.

31705 **RETURN VALUE**

31706 Upon successful completion, *pthread\_attr\_getstackaddr()* and *pthread\_attr\_setstackaddr()* shall  
31707 return a value of 0; otherwise, an error number shall be returned to indicate the error.

31708 The *pthread\_attr\_getstackaddr()* function stores the *stackaddr* attribute value in *stackaddr* if  
31709 successful.

31710 **ERRORS**

31711 No errors are defined.

31712 These functions shall not return an error code of [EINTR].

31713 **EXAMPLES**

31714 None.

31715 **APPLICATION USAGE**

31716 None.

31717 **RATIONALE**

31718 None.

31719 **FUTURE DIRECTIONS**

31720 None.

31721 **SEE ALSO**

31722 *pthread\_attr\_destroy()*, *pthread\_attr\_getdetachstate()*, *pthread\_attr\_getstacksize()*, *pthread\_create()*,  
31723 the Base Definitions volume of IEEE Std. 1003.1-200x, <limits.h>, <pthread.h>

31724 **CHANGE HISTORY**

31725 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31726 **Issue 6**

31727 The *pthread\_attr\_getstackaddr()* and *pthread\_attr\_setstackaddr()* functions are marked as part of  
31728 the Threads and Thread Stack Address Attribute options.

31729 The **restrict** keyword is added to the *pthread\_attr\_getstackaddr()* prototype for alignment with the  
31730 ISO/IEC 9899:1999 standard.

31731 **NAME**

31732 pthread\_attr\_getstacksize, pthread\_attr\_setstacksize — get and set stacksize attribute

31733 **SYNOPSIS**

31734 THR TSA #include &lt;pthread.h&gt;

```
31735 int pthread_attr_getstacksize(const pthread_attr_t *restrict attr,
31736 size_t *restrict stacksize);
31737 int pthread_attr_setstacksize(pthread_attr_t *attr, size_t stacksize);
31738
```

31739 **DESCRIPTION**31740 The *pthread\_attr\_getstacksize()* and *pthread\_attr\_setstacksize()* functions, respectively, get and set  
31741 the thread creation *stacksize* attribute in the *attr* object.31742 The *stacksize* attribute defines the minimum stack size (in bytes) allocated for the created threads  
31743 stack.31744 **RETURN VALUE**31745 Upon successful completion, *pthread\_attr\_getstacksize()* and *pthread\_attr\_setstacksize()* shall  
31746 return a value of 0; otherwise, an error number shall be returned to indicate the error.31747 The *pthread\_attr\_getstacksize()* function stores the *stacksize* attribute value in *stacksize* if  
31748 successful.31749 **ERRORS**31750 The *pthread\_attr\_setstacksize()* function shall fail if:31751 [EINVAL] The value of *stacksize* is less than {PTHREAD\_STACK\_MIN} or exceeds a  
31752 system-imposed limit.

31753 These functions shall not return an error code of [EINTR].

31754 **EXAMPLES**

31755 None.

31756 **APPLICATION USAGE**

31757 None.

31758 **RATIONALE**

31759 None.

31760 **FUTURE DIRECTIONS**

31761 None.

31762 **SEE ALSO**31763 *pthread\_attr\_destroy()*, *pthread\_attr\_getstackaddr()*, *pthread\_attr\_getdetachstate()*, *pthread\_create()*,  
31764 the Base Definitions volume of IEEE Std. 1003.1-200x, <limits.h>, <pthread.h>31765 **CHANGE HISTORY**

31766 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

31767 **Issue 6**31768 The *pthread\_attr\_getstacksize()* and *pthread\_attr\_setstacksize()* functions are marked as part of the  
31769 Threads and Thread Stack Address Attribute options.31770 The **restrict** keyword is added to the *pthread\_attr\_getstacksize()* prototype for alignment with the  
31771 ISO/IEC 9899:1999 standard.

31772 **NAME**

31773 pthread\_attr\_init — initialize threads attributes object

31774 **SYNOPSIS**

31775 THR #include &lt;pthread.h&gt;

31776 int pthread\_attr\_init(pthread\_attr\_t \*attr);

31777

31778 **DESCRIPTION**31779 Refer to *pthread\_attr\_destroy()*.

31780 **NAME**

31781 pthread\_attr\_setdetachstate — set detachstate attribute

31782 **SYNOPSIS**

31783 THR #include <pthread.h>

31784 int pthread\_attr\_setdetachstate(pthread\_attr\_t \*attr, int detachstate);

31785

31786 **DESCRIPTION**

31787 Refer to *pthread\_attr\_getdetachstate()*.

31788 **NAME**

31789 pthread\_attr\_setguardsize — set thread guardsize attribute

31790 **SYNOPSIS**

31791 xSI #include &lt;pthread.h&gt;

31792 int pthread\_attr\_setguardsize(pthread\_attr\_t \*attr,  
31793 size\_t guardsize);

31794

31795 **DESCRIPTION**31796 Refer to *pthread\_attr\_getguardsize()*.

31797 **NAME**

31798 pthread\_attr\_setinheritsched — set inheritsched attribute (**REALTIME THREADS**)

31799 **SYNOPSIS**

31800 TPS #include <pthread.h>

```
31801 int pthread_attr_setinheritsched(pthread_attr_t *attr,
31802 int inheritsched);
```

31803

31804 **DESCRIPTION**

31805 Refer to *pthread\_attr\_getinheritsched()*.

31806 **NAME**

31807 pthread\_attr\_setschedparam — set schedparam attribute

31808 **SYNOPSIS**

31809 THR #include &lt;pthread.h&gt;

31810 int pthread\_attr\_setschedparam(pthread\_attr\_t \*restrict attr,  
31811 const struct sched\_param \*restrict param);

31812

31813 **DESCRIPTION**31814 Refer to *pthread\_attr\_getschedparam()*.

31815 **NAME**

31816 pthread\_attr\_setschedpolicy — set schedpolicy attribute (**REALTIME THREADS**)

31817 **SYNOPSIS**

31818 TPS #include <pthread.h>

31819 int pthread\_attr\_setschedpolicy(pthread\_attr\_t \*attr, int policy);

31820

31821 **DESCRIPTION**

31822 Refer to *pthread\_attr\_getschedpolicy()*.

31823 **NAME**

31824 pthread\_attr\_setscope — set contentionscope attribute (**REALTIME THREADS**)

31825 **SYNOPSIS**

31826 TPS #include <pthread.h>

31827 int pthread\_attr\_setscope(pthread\_attr\_t \*attr, int contentionscope);

31828

31829 **DESCRIPTION**

31830 Refer to *pthread\_attr\_getscope()*.

31831 **NAME**

31832 pthread\_attr\_setstack — set stack attribute

31833 **SYNOPSIS**

31834 XSI #include <pthread.h>

```
31835 int pthread_attr_setstack(pthread_attr_t *attr, void *stackaddr,
31836 size_t stacksize);
```

31837

31838 **DESCRIPTION**

31839 Refer to *pthread\_attr\_getstack()*.

31840 **NAME**

31841 pthread\_attr\_setstackaddr — set stackaddr attribute

31842 **SYNOPSIS**

31843 THR TSA #include &lt;pthread.h&gt;

31844 int pthread\_attr\_setstackaddr(pthread\_attr\_t \*attr, void \*stackaddr);

31845

31846 **DESCRIPTION**31847 Refer to *pthread\_attr\_getstackaddr()*.

31848 **NAME**

31849 pthread\_attr\_setstacksize — set stacksize attribute

31850 **SYNOPSIS**

31851 THR TSA #include <pthread.h>

31852 int pthread\_attr\_setstacksize(pthread\_attr\_t \*attr, size\_t stacksize);

31853

31854 **DESCRIPTION**

31855 Refer to *pthread\_attr\_getstacksize()*.

31856 **NAME**

31857 pthread\_barrier\_destroy, pthread\_barrier\_init — destroy and initialize a barrier object

31858 **SYNOPSIS**31859 **BAR** #include <pthread.h>

```

31860 int pthread_barrier_destroy(pthread_barrier_t *barrier);
31861 int pthread_barrier_init(pthread_barrier_t *restrict barrier,
31862 const pthread_barrierattr_t *restrict attr, unsigned count);
31863

```

31864 **DESCRIPTION**

31865 The *pthread\_barrier\_destroy()* function destroys the barrier referenced by *barrier* and releases any  
 31866 resources used by the barrier. The effect of subsequent use of the barrier is undefined until the  
 31867 barrier is reinitialized by another call to *pthread\_barrier\_init()*. An implementation may use this  
 31868 function to set *barrier* to an invalid value. The results are undefined if *pthread\_barrier\_destroy()* is  
 31869 called when any thread is blocked on the barrier, or if this function is called with an uninitialized  
 31870 barrier.

31871 The *pthread\_barrier\_init()* function shall allocate any resources required to use the barrier  
 31872 referenced by *barrier* and initializes the barrier with attributes referenced by *attr*. If *attr* is NULL,  
 31873 the default barrier attributes are used; the effect is the same as passing the address of a default  
 31874 barrier attributes object. The results are undefined if *pthread\_barrier\_init()* is called when any  
 31875 thread is blocked on the barrier (that is, has not returned from the *pthread\_barrier\_wait()* call).  
 31876 The results are undefined if a barrier is used without first being initialized. The results are  
 31877 undefined if *pthread\_barrier\_init()* is called specifying an already initialized barrier.

31878 The *count* argument specifies the number of threads that must call *pthread\_barrier\_wait()* before  
 31879 any of them successfully return from the call. The value specified by *count* must be greater than  
 31880 zero.

31881 If the *pthread\_barrier\_init()* function fails, the barrier is not initialized and the contents of *barrier*  
 31882 are undefined.

31883 Only the object referenced by *barrier* may be used for performing synchronization. The result of  
 31884 referring to copies of that object in calls to *pthread\_barrier\_destroy()* or *pthread\_barrier\_wait()* is  
 31885 undefined.

31886 **RETURN VALUE**

31887 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
 31888 be returned to indicate the error.

31889 **ERRORS**

31890 The *pthread\_barrier\_destroy()* function may fail if:

31891 [EBUSY] The implementation has detected an attempt to destroy a barrier while it is in  
 31892 use (for example, while being used in a *pthread\_barrier\_wait()* call) by another  
 31893 thread.

31894 [EINVAL] The value specified by *barrier* is invalid.

31895 The *pthread\_barrier\_init()* function shall fail if:

31896 [EAGAIN] The system lacks the necessary resources to initialize another barrier.

31897 [EINVAL] The value specified by *count* is equal to zero.

31898 [ENOMEM] Insufficient memory exists to initialize the barrier.

- 31899 The *pthread\_barrier\_init()* function may fail if:
- 31900 [EBUSY] The implementation has detected an attempt to reinitialize a barrier while it is  
31901 in use (for example, while being used in a *pthread\_barrier\_wait()* call) by  
31902 another thread.
- 31903 [EINVAL] The value specified by *attr* is invalid.
- 31904 These functions shall not return an error code of [EINTR].
- 31905 **EXAMPLES**
- 31906 None.
- 31907 **APPLICATION USAGE**
- 31908 The *pthread\_barrier\_destroy()* and *pthread\_barrier\_init()* functions are part of the Barriers option  
31909 and need not be provided on all implementations.
- 31910 **RATIONALE**
- 31911 None.
- 31912 **FUTURE DIRECTIONS**
- 31913 None.
- 31914 **SEE ALSO**
- 31915 *pthread\_barrier\_wait()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**pthread.h**>
- 31916 **CHANGE HISTORY**
- 31917 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

31918 **NAME**

31919 pthread\_barrier\_init — initialize a barrier object

31920 **SYNOPSIS**

31921 BAR #include &lt;pthread.h&gt;

31922 int pthread\_barrier\_init(pthread\_barrier\_t \*restrict *barrier*,  
31923 const pthread\_barrierattr\_t \*restrict *attr*, unsigned *count*);

31924

31925 **DESCRIPTION**31926 Refer to *pthread\_barrier\_destroy()*.

31927 **NAME**

31928 pthread\_barrier\_wait — synchronize at a barrier

31929 **SYNOPSIS**

31930 BAR #include &lt;pthread.h&gt;

31931 int pthread\_barrier\_wait(pthread\_barrier\_t \*barrier);

31932

31933 **DESCRIPTION**

31934 The *pthread\_barrier\_wait()* function synchronizes participating threads at the barrier referenced  
31935 by *barrier*. The calling thread blocks (that is, does not return from the *pthread\_barrier\_wait()* call)  
31936 until the required number of threads have called *pthread\_barrier\_wait()* specifying the barrier.

31937 When the required number of threads have called *pthread\_barrier\_wait()* specifying the barrier,  
31938 the constant PTHREAD\_BARRIER\_SERIAL\_THREAD is returned to one unspecified thread  
31939 and zero is returned to each of the remaining threads. At this point, the barrier is reset to the  
31940 state it had as a result of the most recent *pthread\_barrier\_init()* function that referenced it.

31941 The constant PTHREAD\_BARRIER\_SERIAL\_THREAD is defined in <pthread.h> and its value  
31942 is distinct from any other value returned by *pthread\_barrier\_wait()*.

31943 The results are undefined if this function is called with an uninitialized barrier.

31944 If a signal is delivered to a thread blocked on a barrier, upon return from the signal handler the  
31945 thread shall resume waiting at the barrier if the barrier wait has not completed (that is, if the  
31946 required number of threads have not arrived at the barrier during the execution of the signal  
31947 handler); otherwise, the thread shall continue as normal from the completed barrier wait. Until  
31948 the thread in the signal handler returns from it, it is unspecified whether other threads may  
31949 proceed past the barrier once they have all reached it.

31950 A thread that has blocked on a barrier shall not prevent any unblocked thread that is eligible to  
31951 use the same processing resources from eventually making forward progress in its execution.  
31952 Eligibility for processing resources shall be determined by the scheduling policy.

31953 **RETURN VALUE**

31954 Upon successful completion, the *pthread\_barrier\_wait()* function shall return  
31955 PTHREAD\_BARRIER\_SERIAL\_THREAD for a single (arbitrary) thread synchronized at the  
31956 barrier and zero for each of the other threads. Otherwise, an error number shall be returned to  
31957 indicate the error.

31958 **ERRORS**

31959 The *pthread\_barrier\_wait()* function may fail if:

31960 [EINVAL] The value specified by *barrier* does not refer to an initialized barrier object.

31961 This function shall not return an error code of [EINTR].

31962 **EXAMPLES**

31963 None.

31964 **APPLICATION USAGE**

31965 Applications using this function may be subject to priority inversion, as discussed in the Base  
31966 Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

31967 The *pthread\_barrier\_wait()* function is part of the Barriers option and need not be provided on all  
31968 implementations.

31969 **RATIONALE**

31970 None.

31971 **FUTURE DIRECTIONS**

31972 None.

31973 **SEE ALSO**31974 *pthread\_barrier\_destroy()*, *pthread\_barrier\_init()*, the Base Definitions volume of |

31975 IEEE Std. 1003.1-200x, &lt;pthread.h&gt; |

31976 **CHANGE HISTORY**

31977 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

31978 In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

31979 **NAME**

31980 pthread\_barrierattr\_destroy, pthread\_barrierattr\_init — destroy and initialize barrier attributes  
31981 object

31982 **SYNOPSIS**

```
31983 BAR #include <pthread.h>
```

```
31984 int pthread_barrierattr_destroy(pthread_barrierattr_t *attr);
```

```
31985 int pthread_barrierattr_init(pthread_barrierattr_t *attr);
```

```
31986
```

31987 **DESCRIPTION**

31988 The *pthread\_barrierattr\_destroy()* function destroys a barrier attributes object. The effect of  
31989 subsequent use of the object is undefined until the object is reinitialized by another call to  
31990 *pthread\_barrierattr\_init()*. An implementation may cause *pthread\_barrierattr\_destroy()* to set the  
31991 object referenced by *attr* to an invalid value.

31992 The *pthread\_barrierattr\_init()* function initializes a barrier attributes object *attr* with the default  
31993 value for all of the attributes defined by the implementation.

31994 The results are undefined if *pthread\_barrierattr\_init()* is called specifying an already initialized  
31995 barrier attributes object.

31996 After a barrier attributes object has been used to initialize one or more barriers, any function  
31997 affecting the attributes object (including destruction) does not affect any previously initialized  
31998 barrier.

31999 **RETURN VALUE**

32000 If successful, the *pthread\_barrierattr\_destroy()* and *pthread\_barrierattr\_init()* functions shall return  
32001 zero; otherwise, an error number shall be returned to indicate the error.

32002 **ERRORS**

32003 The *pthread\_barrierattr\_destroy()* function may fail if:

32004 [EINVAL] The value specified by *attr* is invalid.

32005 The *pthread\_barrierattr\_init()* function shall fail if:

32006 [ENOMEM] Insufficient memory exists to initialize the barrier attributes object.

32007 These functions shall not return an error code of [EINTR].

32008 **EXAMPLES**

32009 None.

32010 **APPLICATION USAGE**

32011 The *pthread\_barrierattr\_destroy()* and *pthread\_barrierattr\_init()* functions are part of the Barriers  
32012 option and need not be provided on all implementations.

32013 **RATIONALE**

32014 None.

32015 **FUTURE DIRECTIONS**

32016 None.

32017 **SEE ALSO**

32018 *pthread\_barrierattr\_getshared()*, *pthread\_barrierattr\_setshared()*, the Base Definitions volume of  
32019 IEEE Std. 1003.1-200x, <pthread.h>.

32020 **CHANGE HISTORY**

32021 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

32022 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

32023 **NAME**

32024 pthread\_barrierattr\_getpshared, pthread\_barrierattr\_setpshared — get and set process-shared  
 32025 attribute of barrier attributes object

32026 **SYNOPSIS**

32027 BAR TSH #include <pthread.h>

```
32028 int pthread_barrierattr_getpshared(const pthread_barrierattr_t *restrict attr, |
32029 int *restrict pshared); |
32030 int pthread_barrierattr_setpshared(pthread_barrierattr_t *attr, |
32031 int pshared); |
32032
```

32033 **DESCRIPTION**

32034 The process-shared attribute is set to PTHREAD\_PROCESS\_SHARED to permit a barrier to be  
 32035 operated upon by any thread that has access to the memory where the barrier is allocated. If the  
 32036 process-shared attribute is PTHREAD\_PROCESS\_PRIVATE, the barrier shall only be operated  
 32037 upon by threads created within the same process as the thread that initialized the barrier; if  
 32038 threads of different processes attempt to operate on such a barrier, the behavior is undefined.  
 32039 The default value of the attribute shall be PTHREAD\_PROCESS\_PRIVATE. Both constants  
 32040 PTHREAD\_PROCESS\_SHARED and PTHREAD\_PROCESS\_PRIVATE are defined in  
 32041 <pthread.h>.

32042 The *pthread\_barrierattr\_getpshared()* function obtains the value of the process-shared attribute  
 32043 from the attributes object referenced by *attr*. The *pthread\_barrierattr\_setpshared()* function is used  
 32044 to set the process-shared attribute in an initialized attributes object referenced by *attr*.

32045 Additional attributes, their default values, and the names of the associated functions to get and  
 32046 set those attribute values are implementation-defined.

32047 **RETURN VALUE**

32048 If successful, the *pthread\_barrierattr\_getpshared()* function shall return zero and store the value of  
 32049 the process-shared attribute of *attr* into the object referenced by the *pshared* parameter.  
 32050 Otherwise, an error number shall be returned to indicate the error.

32051 If successful, the *pthread\_barrierattr\_setpshared()* function shall return zero; otherwise, an error  
 32052 number shall be returned to indicate the error.

32053 **ERRORS**

32054 These functions may fail if:

32055 [EINVAL] The value specified by *attr* is invalid.

32056 The *pthread\_barrierattr\_setpshared()* function may fail if:

32057 [EINVAL] The new value specified for the process-shared attribute is not one of the legal  
 32058 values PTHREAD\_PROCESS\_SHARED or PTHREAD\_PROCESS\_PRIVATE.

32059 These functions shall not return an error code of [EINTR].

32060 **EXAMPLES**

32061 None.

32062 **APPLICATION USAGE**

32063 The *pthread\_barrierattr\_getpshared()* and *pthread\_barrierattr\_setpshared()* functions are part of the  
32064 Barriers option and need not be provided on all implementations.

32065 **RATIONALE**

32066 None.

32067 **FUTURE DIRECTIONS**

32068 None.

32069 **SEE ALSO**

32070 *pthread\_barrier\_init()*, *pthread\_barrierattr\_destroy()*, *pthread\_barrierattr\_init()*, the Base Definitions  
32071 volume of IEEE Std. 1003.1-200x, <pthread.h>

32072 **CHANGE HISTORY**

32073 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000

32074 **NAME**

32075 pthread\_barrierattr\_init — initialize barrier attributes object

32076 **SYNOPSIS**

32077 BAR #include <pthread.h>

32078 int pthread\_barrierattr\_init(pthread\_barrierattr\_t \*attr);

32079

32080 **DESCRIPTION**

32081 Refer to *pthread\_barrierattr\_destroy()*.

32082 **NAME**

32083 pthread\_barrierattr\_setpshared — set process-shared attribute of barrier attributes object

32084 **SYNOPSIS**

32085 BAR TSH #include &lt;pthread.h&gt;

32086 int pthread\_barrierattr\_setpshared(pthread\_barrierattr\_t \*attr,  
32087 int pshared);

32088

32089 **DESCRIPTION**32090 Refer to *pthread\_barrierattr\_getpshared()*.

32091 **NAME**

32092 pthread\_cancel — cancel execution of a thread

32093 **SYNOPSIS**

32094 THR #include &lt;pthread.h&gt;

32095 int pthread\_cancel(pthread\_t thread);

32096

32097 **DESCRIPTION**

32098 The *pthread\_cancel()* function requests that *thread* be canceled. The target thread's cancelability state and type determines when the cancellation takes effect. When the cancellation is acted on, 32099 the cancellation cleanup handlers for *thread* are called. When the last cancellation cleanup handler 32100 returns, the thread-specific data destructor functions shall be called for *thread*. When the last 32101 destructor function returns, *thread* shall be terminated. 32102

32103 The cancellation processing in the target thread runs asynchronously with respect to the calling 32104 thread returning from *pthread\_cancel()*.

32105 **RETURN VALUE**

32106 If successful, the *pthread\_cancel()* function shall return zero; otherwise, an error number shall be 32107 returned to indicate the error.

32108 **ERRORS**32109 The *pthread\_cancel()* function may fail if:

32110 [ESRCH] No thread could be found corresponding to that specified by the given thread ID. 32111

32112 The *pthread\_cancel()* function shall not return an error code of [EINTR].

32113 **EXAMPLES**

32114 None.

32115 **APPLICATION USAGE**

32116 None.

32117 **RATIONALE**

32118 Two alternative functions were considered to sending the cancellation notification to a thread. 32119 One would be to define a new SIGCANCEL signal that had the cancellation semantics when 32120 delivered; the other was to define the new *pthread\_cancel()* function, which would trigger the 32121 cancellation semantics.

32122 The advantage of a new signal was that so much of the delivery criteria were identical to that 32123 used when trying to deliver a signal that making cancellation notification a signal was seen as 32124 consistent. Indeed, many implementations implement cancellation using a special signal. On the 32125 other hand, there would be no signal functions that could be used with this signal except 32126 *pthread\_kill()*, and the behavior of the delivered cancellation signal would be unlike any 32127 previously existing defined signal.

32128 The benefits of a special function include the recognition that this signal would be defined 32129 because of the similar delivery criteria and that this is the only common behavior between a 32130 cancellation request and a signal. In addition, the cancellation delivery mechanism does not have 32131 to be implemented as a signal. There are also strong, if not stronger, parallels with language 32132 exception mechanisms than with signals that are potentially obscured if the delivery mechanism 32133 is visibly closer to signals.

32134 In the end, it was considered that as there were so many exceptions to the use of the new signal 32135 with existing signals functions that it would be misleading. A special function has resolved this

32136 problem. This function was carefully defined so that an implementation wishing to provide the  
32137 cancelation functions on top of signals could do so. The special function also means that  
32138 implementations are not obliged to implement cancelation with signals.

32139 **FUTURE DIRECTIONS**

32140 None.

32141 **SEE ALSO**

32142 *pthread\_exit()*, *pthread\_cond\_wait()*, *pthread\_cond\_timedwait()*, *pthread\_join()*,  
32143 *pthread\_setcancelstate()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**pthread.h**>

32144 **CHANGE HISTORY**

32145 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32146 **Issue 6**

32147 The *pthread\_cancel()* function is marked as part of the Threads option.

32148 **NAME**

32149 pthread\_cleanup\_pop, pthread\_cleanup\_push — establish cancelation handlers

32150 **SYNOPSIS**

32151 THR #include &lt;pthread.h&gt;

32152 void pthread\_cleanup\_pop(int *execute*);32153 void pthread\_cleanup\_push(void (\**routine*)(void\*), void \**arg*);

32154

32155 **DESCRIPTION**32156 The *pthread\_cleanup\_pop()* function shall remove the routine at the top of the calling thread's  
32157 cancelation cleanup stack and optionally invoke it (if *execute* is non-zero).32158 The *pthread\_cleanup\_push()* function shall push the specified cancelation cleanup handler *routine*  
32159 onto the calling thread's cancelation cleanup stack. The cancelation cleanup handler shall be  
32160 popped from the cancelation cleanup stack and invoked with the argument *arg* when:

- 32161 • The thread exits (that is, calls *pthread\_exit()*).
- 32162 • The thread acts upon a cancelation request.
- 32163 • The thread calls *pthread\_cleanup\_pop()* with a non-zero *execute* argument.

32164 These functions may be implemented as macros. The application shall ensure that they appear  
32165 as statements, and in pairs within the same lexical scope (that is, the *pthread\_cleanup\_push()*  
32166 macro may be thought to expand to a token list whose first token is '{' with  
32167 *pthread\_cleanup\_pop()* expanding to a token list whose last token is the corresponding '}').32168 The effect of calling *longjmp()* or *siglongjmp()* is undefined if there have been any calls to  
32169 *pthread\_cleanup\_push()* or *pthread\_cleanup\_pop()* made without the matching call since the jump  
32170 buffer was filled. The effect of calling *longjmp()* or *siglongjmp()* from inside a cancelation  
32171 cleanup handler is also undefined unless the jump buffer was also filled in the cancelation  
32172 cleanup handler.32173 **RETURN VALUE**32174 The *pthread\_cleanup\_push()* and *pthread\_cleanup\_pop()* functions shall return no value.32175 **ERRORS**

32176 No errors are defined.

32177 These functions shall not return an error code of [EINTR].

32178 **EXAMPLES**32179 The following is an example using thread primitives to implement a cancelable, writers-priority  
32180 read-write lock:

```

32181 typedef struct {
32182 pthread_mutex_t lock;
32183 pthread_cond_t rcond,
32184 wcond;
32185 int lock_count; /* < 0 .. Held by writer. */
32186 /* > 0 .. Held by lock_count readers. */
32187 /* = 0 .. Held by nobody. */
32188 int waiting_writers; /* Count of waiting writers. */
32189 } rwlock;
32190
32191 void
32192 waiting_reader_cleanup(void *arg)
32193 {

```

```

32193 rwlock *l;
32194 l = (rwlock *) arg;
32195 pthread_mutex_unlock(&l->lock);
32196 }
32197 void
32198 lock_for_read(rwlock *l)
32199 {
32200 pthread_mutex_lock(&l->lock);
32201 pthread_cleanup_push(waiting_reader_cleanup, l);
32202 while ((l->lock_count < 0) && (l->waiting_writers != 0))
32203 pthread_cond_wait(&l->rcond, &l->lock);
32204 l->lock_count++;
32205 /*
32206 * Note the pthread_cleanup_pop executes
32207 * waiting_reader_cleanup.
32208 */
32209 pthread_cleanup_pop(1);
32210 }
32211 void
32212 release_read_lock(rwlock *l)
32213 {
32214 pthread_mutex_lock(&l->lock);
32215 if (--l->lock_count == 0)
32216 pthread_cond_signal(&l->wcond);
32217 pthread_mutex_unlock(l);
32218 }
32219 void
32220 waiting_writer_cleanup(void *arg)
32221 {
32222 rwlock *l;
32223 l = (rwlock *) arg;
32224 if ((--l->waiting_writers == 0) && (l->lock_count >= 0)) {
32225 /*
32226 * This only happens if we have been canceled.
32227 */
32228 pthread_cond_broadcast(&l->wcond);
32229 }
32230 pthread_mutex_unlock(&l->lock);
32231 }
32232 void
32233 lock_for_write(rwlock *l)
32234 {
32235 pthread_mutex_lock(&l->lock);
32236 l->waiting_writers++;
32237 pthread_cleanup_push(waiting_writer_cleanup, l);
32238 while (l->lock_count != 0)
32239 pthread_cond_wait(&l->wcond, &l->lock);
32240 l->lock_count = -1;
32241 /*

```

```
32242 * Note the pthread_cleanup_pop executes
32243 * waiting_writer_cleanup.
32244 */
32245 pthread_cleanup_pop(1);
32246 }

32247 void
32248 release_write_lock(rwlock *l)
32249 {
32250 pthread_mutex_lock(&l->lock);
32251 l->lock_count = 0;
32252 if (l->waiting_writers == 0)
32253 pthread_cond_broadcast(&l->rcond)
32254 else
32255 pthread_cond_signal(&l->wcond);
32256 pthread_mutex_unlock(&l->lock);
32257 }

32258 /*
32259 * This function is called to initialize the read/write lock.
32260 */
32261 void
32262 initialize_rwlock(rwlock *l)
32263 {
32264 pthread_mutex_init(&l->lock, pthread_mutexattr_default);
32265 pthread_cond_init(&l->wcond, pthread_condattr_default);
32266 pthread_cond_init(&l->rcond, pthread_condattr_default);
32267 l->lock_count = 0;
32268 l->waiting_writers = 0;
32269 }

32270 reader_thread()
32271 {
32272 lock_for_read(&lock);
32273 pthread_cleanup_push(release_read_lock, &lock);
32274 /*
32275 * Thread has read lock.
32276 */
32277 pthread_cleanup_pop(1);
32278 }

32279 writer_thread()
32280 {
32281 lock_for_write(&lock);
32282 pthread_cleanup_push(release_write_lock, &lock);
32283 /*
32284 * Thread has write lock.
32285 */
32286 pthread_cleanup_pop(1);
32287 }
```

32288 **APPLICATION USAGE**

32289 The two routines that push and pop cancelation cleanup handlers, *pthread\_cleanup\_push()* and  
 32290 *pthread\_cleanup\_pop()*, can be thought of as left and right parentheses. They always need to be  
 32291 matched.

32292 **RATIONALE**

32293 The restriction that the two routines that push and pop cancelation cleanup handlers,  
 32294 *pthread\_cleanup\_push()* and *pthread\_cleanup\_pop()*, have to appear in the same lexical scope  
 32295 allows for efficient macro or compiler implementations and efficient storage management. A  
 32296 sample implementation of these routines as macros might look like this:

```
32297 #define pthread_cleanup_push(rtn,arg) { \
32298 struct _pthread_handler_rec __cleanup_handler, **__head; \
32299 __cleanup_handler.rtn = rtn; \
32300 __cleanup_handler.arg = arg; \
32301 (void) pthread_getspecific(_pthread_handler_key, &__head); \
32302 __cleanup_handler.next = *__head; \
32303 *__head = &__cleanup_handler;
32304
32305 #define pthread_cleanup_pop(ex) \
32306 *__head = __cleanup_handler.next; \
32307 if (ex) (*__cleanup_handler.rtn)(__cleanup_handler.arg); \
32308 }
```

32308 A more ambitious implementation of these routines might do even better by allowing the  
 32309 compiler to note that the cancelation cleanup handler is a constant and can be expanded inline.

32310 This volume of IEEE Std. 1003.1-200x currently leaves unspecified the effect of calling *longjmp()*  
 32311 from a signal handler executing in a POSIX System Interfaces function. If an implementation  
 32312 wants to allow this and give the programmer reasonable behavior, the *longjmp()* function has to  
 32313 call all cancelation cleanup handlers that have been pushed but not popped since the time  
 32314 *setjmp()* was called.

32315 Consider a multi-threaded function called by a thread that uses signals. If a signal were  
 32316 delivered to a signal handler during the operation of *qsort()* and that handler were to call  
 32317 *longjmp()* (which, in turn, did *not* call the cancelation cleanup handlers) the helper threads  
 32318 created by the *qsort()* function would not be canceled. Instead, they would continue to execute  
 32319 and write into the argument array even though the array might have been popped off of the  
 32320 stack.

32321 Note that the specified cleanup handling mechanism is especially tied to the C language and,  
 32322 while the requirement for a uniform mechanism for expressing cleanup is language-  
 32323 independent, the mechanism used in other languages may be quite different. In addition, this  
 32324 mechanism is really only necessary due to the lack of a real exception mechanism in the C  
 32325 language, which would be the ideal solution.

32326 There is no notion of a cancelation cleanup-safe function. If an application has no cancelation  
 32327 points in its signal handlers, blocks any signal whose handler may have cancelation points while  
 32328 calling async-unsafe functions, or disables cancelation while calling async-unsafe functions, all  
 32329 functions may be safely called from cancelation cleanup routines.

32330 **FUTURE DIRECTIONS**

32331 None.

32332 **SEE ALSO**

32333        *pthread\_cancel()*, *pthread\_setcancelstate()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
32334        <pthread.h>

32335 **CHANGE HISTORY**

32336        First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32337 **Issue 6**

32338        The *pthread\_cleanup\_pop()* and *pthread\_cleanup\_push()* functions are marked as part of the  
32339        Threads option.

32340        The APPLICATION USAGE section is added.

32341        The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

32342 **NAME**

32343 pthread\_cond\_broadcast, pthread\_cond\_signal — broadcast or signal a condition

32344 **SYNOPSIS**

32345 THR #include &lt;pthread.h&gt;

32346 int pthread\_cond\_broadcast(pthread\_cond\_t \*cond);

32347 int pthread\_cond\_signal(pthread\_cond\_t \*cond);

32348

32349 **DESCRIPTION**

32350 These functions are used to unblock threads blocked on a condition variable.

32351 The *pthread\_cond\_broadcast()* function shall unblock all threads currently blocked on the  
32352 specified condition variable *cond*.32353 The *pthread\_cond\_signal()* function shall unblock at least one of the threads that are blocked on  
32354 the specified condition variable *cond* (if any threads are blocked on *cond*).32355 If more than one thread is blocked on a condition variable, the scheduling policy determines the  
32356 order in which threads are unblocked. When each thread unblocked as a result of a  
32357 *pthread\_cond\_broadcast()* or *pthread\_cond\_signal()* returns from its call to *pthread\_cond\_wait()* or  
32358 *pthread\_cond\_timedwait()*, the thread owns the mutex with which it called *pthread\_cond\_wait()* or  
32359 *pthread\_cond\_timedwait()*. The thread(s) that are unblocked shall contend for the mutex  
32360 according to the scheduling policy (if applicable), and as if each had called *pthread\_mutex\_lock()*.32361 The *pthread\_cond\_broadcast()* or *pthread\_cond\_signal()* functions may be called by a thread  
32362 whether or not it currently owns the mutex that threads calling *pthread\_cond\_wait()* or  
32363 *pthread\_cond\_timedwait()* have associated with the condition variable during their waits;  
32364 however, if predictable scheduling behavior is required, then that mutex shall be locked by the  
32365 thread calling *pthread\_cond\_broadcast()* or *pthread\_cond\_signal()*.32366 The *pthread\_cond\_broadcast()* and *pthread\_cond\_signal()* functions have no effect if there are no  
32367 threads currently blocked on *cond*.32368 **RETURN VALUE**32369 If successful, the *pthread\_cond\_broadcast()* and *pthread\_cond\_signal()* functions shall return zero;  
32370 otherwise, an error number shall be returned to indicate the error.32371 **ERRORS**32372 The *pthread\_cond\_broadcast()* and *pthread\_cond\_signal()* function may fail if:32373 [EINVAL] The value *cond* does not refer to an initialized condition variable.

32374 These functions shall not return an error code of [EINTR].

32375 **EXAMPLES**

32376 None.

32377 **APPLICATION USAGE**32378 The *pthread\_cond\_broadcast()* function is used whenever the shared-variable state has been  
32379 changed in a way that more than one thread can proceed with its task. Consider a single  
32380 producer/multiple consumer problem, where the producer can insert multiple items on a list  
32381 that is accessed one item at a time by the consumers. By calling the *pthread\_cond\_broadcast()*  
32382 function, the producer would notify all consumers that might be waiting, and thereby the  
32383 application would receive more throughput on a multiprocessor. In addition,  
32384 *pthread\_cond\_broadcast()* makes it easier to implement a read-write lock. The  
32385 *pthread\_cond\_broadcast()* function is needed in order to wake up all waiting readers when a  
32386 writer releases its lock. Finally, the two-phase commit algorithm can use this broadcast function

32387 to notify all clients of an impending transaction commit.

32388 It is not safe to use the *pthread\_cond\_signal()* function in a signal handler that is invoked  
32389 asynchronously. Even if it were safe, there would still be a race between the test of the Boolean  
32390 *pthread\_cond\_wait()* that could not be efficiently eliminated.

32391 Mutexes and condition variables are thus not suitable for releasing a waiting thread by signaling  
32392 from code running in a signal handler.

### 32393 RATIONALE

#### 32394 Multiple Awakenings by Condition Signal

32395 On a multiprocessor, it may be impossible for an implementation of *pthread\_cond\_signal()* to  
32396 avoid the unblocking of more than one thread blocked on a condition variable. For example,  
32397 consider the following partial implementation of *pthread\_cond\_wait()* and *pthread\_cond\_signal()*,  
32398 executed by two threads in the order given. One thread is trying to wait on the condition  
32399 variable, another is concurrently executing *pthread\_cond\_signal()*, while a third thread is already  
32400 waiting.

```
32401 pthread_cond_wait(mutex, cond):
32402 value = cond->value; /* 1 */
32403 pthread_mutex_unlock(mutex); /* 9 */
32404 pthread_mutex_lock(cond->mutex); /* 10 */
32405 if (value == cond->value) { /* 11 */
32406 me->next_cond = cond->waiter;
32407 cond->waiter = me;
32408 pthread_mutex_unlock(cond->mutex);
32409 unable_to_run(me);
32410 } else
32411 pthread_mutex_unlock(cond->mutex); /* 12 */
32412 pthread_mutex_lock(mutex); /* 13 */

32413 pthread_cond_signal(cond):
32414 pthread_mutex_lock(cond->mutex); /* 2 */
32415 cond->value++; /* 3 */
32416 if (cond->waiter) { /* 4 */
32417 sleeper = cond->waiter; /* 5 */
32418 cond->waiter = sleeper->next_cond; /* 6 */
32419 able_to_run(sleeper); /* 7 */
32420 }
32421 pthread_mutex_unlock(cond->mutex); /* 8 */
```

32422 The effect is that more than one thread can return from its call to *pthread\_cond\_wait()* or  
32423 *pthread\_cond\_timedwait()* as a result of one call to *pthread\_cond\_signal()*. This effect is called  
32424 “spurious wakeup”. Note that the situation is self-correcting in that the number of threads that  
32425 are so awakened is finite; for example, the next thread to call *pthread\_cond\_wait()* after the  
32426 sequence of events above blocks.

32427 While this problem could be resolved, the loss of efficiency for a fringe condition that occurs  
32428 only rarely is unacceptable, especially given that one has to check the predicate associated with a  
32429 condition variable anyway. Correcting this problem would unnecessarily reduce the degree of  
32430 concurrency in this basic building block for all higher-level synchronization operations.

32431 An added benefit of allowing spurious wakeups is that applications are forced to code a  
32432 predicate-testing-loop around the condition wait. This also makes the application tolerate  
32433 superfluous condition broadcasts or signals on the same condition variable that may be coded in

32434 some other part of the application. The resulting applications are thus more robust. Therefore,  
32435 IEEE Std. 1003.1-200x explicitly documents that spurious wakeups may occur.

32436 **FUTURE DIRECTIONS**

32437 None.

32438 **SEE ALSO**

32439 *pthread\_cond\_destroy()*, *pthread\_cond\_timedwait()*, *pthread\_cond\_wait()*, the Base Definitions  
32440 volume of IEEE Std. 1003.1-200x, <pthread.h>

32441 **CHANGE HISTORY**

32442 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32443 **Issue 6**

32444 The *pthread\_cond\_broadcast()* and *pthread\_cond\_signal()* functions are marked as part of the  
32445 Threads option.

32446 The APPLICATION USAGE section is added.

## 32447 NAME

32448 pthread\_cond\_destroy, pthread\_cond\_init — destroy and initialize condition variables

## 32449 SYNOPSIS

32450 THR #include &lt;pthread.h&gt;

```

32451 int pthread_cond_destroy(pthread_cond_t *cond);
32452 int pthread_cond_init(pthread_cond_t *restrict cond,
32453 const pthread_condattr_t *restrict attr);
32454 pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
32455

```

## 32456 DESCRIPTION

32457 The *pthread\_cond\_destroy()* function destroys the given condition variable specified by *cond*; the  
 32458 object becomes, in effect, uninitialized. An implementation may cause *pthread\_cond\_destroy()* to  
 32459 set the object referenced by *cond* to an invalid value. A destroyed condition variable object can be  
 32460 re-initialized using *pthread\_cond\_init()*; the results of otherwise referencing the object after it has  
 32461 been destroyed are undefined.

32462 It shall be safe to destroy an initialized condition variable upon which no threads are currently  
 32463 blocked. Attempting to destroy a condition variable upon which other threads are currently  
 32464 blocked results in undefined behavior.

32465 The *pthread\_cond\_init()* function initializes the condition variable referenced by *cond* with  
 32466 attributes referenced by *attr*. If *attr* is NULL, the default condition variable attributes are used;  
 32467 the effect is the same as passing the address of a default condition variable attributes object.  
 32468 Upon successful initialization, the state of the condition variable becomes initialized.

32469 Only *cond* itself may be used for performing synchronization. The result of referring to copies of  
 32470 *cond* in calls to *pthread\_cond\_wait()*, *pthread\_cond\_timedwait()*, *pthread\_cond\_signal()*,  
 32471 *pthread\_cond\_broadcast()*, and *pthread\_cond\_destroy()* is undefined.

32472 Attempting to initialize an already initialized condition variable results in undefined behavior.

32473 In cases where default condition variable attributes are appropriate, the macro  
 32474 PTHREAD\_COND\_INITIALIZER can be used to initialize condition variables that are statically  
 32475 allocated. The effect shall be equivalent to dynamic initialization by a call to *pthread\_cond\_init()*  
 32476 with parameter *attr* specified as NULL, except that no error checks are performed.

## 32477 RETURN VALUE

32478 If successful, the *pthread\_cond\_destroy()* and *pthread\_cond\_init()* functions shall return zero;  
 32479 otherwise, an error number shall be returned to indicate the error.

32480 The [EBUSY] and [EINVAL] error checks, if implemented, shall act as if they were performed  
 32481 immediately at the beginning of processing for the function and caused an error return prior to  
 32482 modifying the state of the condition variable specified by *cond*.

## 32483 ERRORS

32484 The *pthread\_cond\_destroy()* function may fail if:

32485 [EBUSY] The implementation has detected an attempt to destroy the object referenced  
 32486 by *cond* while it is referenced (for example, while being used in a  
 32487 *pthread\_cond\_wait()* or *pthread\_cond\_timedwait()*) by another thread.

32488 [EINVAL] The value specified by *cond* is invalid.

32489 The *pthread\_cond\_init()* function shall fail if:

32490 [EAGAIN] The system lacked the necessary resources (other than memory) to initialize  
 32491 another condition variable.

32492 [ENOMEM] Insufficient memory exists to initialize the condition variable. |

32493 The `pthread_cond_init()` function may fail if:

32494 [EBUSY] The implementation has detected an attempt to re-initialize the object |  
 32495 referenced by `cond`, a previously initialized, but not yet destroyed, condition |  
 32496 variable.

32497 [EINVAL] The value specified by `attr` is invalid. |

32498 These functions shall not return an error code of [EINTR]. |

32499 **EXAMPLES**

32500 A condition variable can be destroyed immediately after all the threads that are blocked on it are  
 32501 awakened. For example, consider the following code:

```
32502 struct list {
32503 pthread_mutex_t lm;
32504 ...
32505 }
32506 struct elt {
32507 key k;
32508 int busy;
32509 pthread_cond_t notbusy;
32510 ...
32511 }
32512 /* Find a list element and reserve it. */
32513 struct elt *
32514 list_find(struct list *lp, key k)
32515 {
32516 struct elt *ep;
32517 pthread_mutex_lock(&lp->lm);
32518 while ((ep = find_elt(l, k) != NULL) && ep->busy)
32519 pthread_cond_wait(&ep->notbusy, &lp->lm);
32520 if (ep != NULL)
32521 ep->busy = 1;
32522 pthread_mutex_unlock(&lp->lm);
32523 return(ep);
32524 }
32525 delete_elt(struct list *lp, struct elt *ep)
32526 {
32527 pthread_mutex_lock(&lp->lm);
32528 assert(ep->busy);
32529 ... remove ep from list ...
32530 ep->busy = 0; /* Paranoid. */
32531 (A) pthread_cond_broadcast(&ep->notbusy);
32532 pthread_mutex_unlock(&lp->lm);
32533 (B) pthread_cond_destroy(&rp->notbusy);
32534 free(ep);
32535 }
```

32536 In this example, the condition variable and its list element may be freed (line B) immediately  
 32537 after all threads waiting for it are awakened (line A), since the mutex and the code ensure that no  
 32538 other thread can touch the element to be deleted.

## 32539 APPLICATION USAGE

32540 None.

## 32541 RATIONALE

32542 See *pthread\_mutex\_init()*; a similar rationale applies to condition variables.

## 32543 FUTURE DIRECTIONS

32544 None.

## 32545 SEE ALSO

32546 *pthread\_cond\_broadcast()*, *pthread\_cond\_signal()*, *pthread\_cond\_timedwait()*, *pthread\_cond\_wait()*,  
32547 the Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

## 32548 CHANGE HISTORY

32549 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

## 32550 Issue 6

32551 The *pthread\_cond\_destroy()* and *pthread\_cond\_init()* functions are marked as part of the Threads  
32552 option.

32553 IEEE PASC Interpretation 1003.1c #34 is applied, updating the DESCRIPTION.

32554 The **restrict** keyword is added to the *pthread\_cond\_init()* prototype for alignment with the  
32555 ISO/IEC 9899:1999 standard.

32556 **NAME**

32557 pthread\_cond\_init — initialize condition variables

32558 **SYNOPSIS**

32559 THR #include &lt;pthread.h&gt;

32560 int pthread\_cond\_init(pthread\_cond\_t \*restrict cond,

32561 const pthread\_condattr\_t \*restrict attr);

32562 pthread\_cond\_t cond = PTHREAD\_COND\_INITIALIZER;

32563

32564 **DESCRIPTION**32565 Refer to *pthread\_cond\_destroy()*.

32566 **NAME**

32567 pthread\_cond\_signal — signal a condition

32568 **SYNOPSIS**

32569 THR #include <pthread.h>

32570 int pthread\_cond\_signal(pthread\_cond\_t \*cond);

32571

32572 **DESCRIPTION**

32573 Refer to *pthread\_cond\_broadcast()*.

32574 **NAME**

32575 pthread\_cond\_timedwait, pthread\_cond\_wait — wait on a condition

32576 **SYNOPSIS**

32577 THR #include &lt;pthread.h&gt;

```

32578 int pthread_cond_timedwait(pthread_cond_t *restrict cond,
32579 pthread_mutex_t *restrict mutex,
32580 const struct timespec *restrict abstime);
32581 int pthread_cond_wait(pthread_cond_t *restrict cond,
32582 pthread_mutex_t *restrict mutex);
32583

```

32584 **DESCRIPTION**

32585 The *pthread\_cond\_timedwait()* and *pthread\_cond\_wait()* functions are used to block on a condition  
 32586 variable. They shall be called with *mutex* locked by the calling thread or undefined behavior  
 32587 results.

32588 These functions atomically release *mutex* and cause the calling thread to block on the condition  
 32589 variable *cond*; atomically here means “atomically with respect to access by another thread to the  
 32590 mutex and then the condition variable”. That is, if another thread is able to acquire the mutex  
 32591 after the about-to-block thread has released it, then a subsequent call to *pthread\_cond\_broadcast()*  
 32592 or *pthread\_cond\_signal()* in that thread shall behave as if it were issued after the about-to-block  
 32593 thread has blocked.

32594 Upon successful return, the mutex has been locked and is owned by the calling thread.

32595 When using condition variables there is always a Boolean predicate involving shared variables  
 32596 associated with each condition wait that is true if the thread should proceed. Spurious wakeups  
 32597 from the *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()* functions may occur. Since the return  
 32598 from *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()* does not imply anything about the value  
 32599 of this predicate, the predicate should be re-evaluated upon such return.

32600 The effect of using more than one mutex for concurrent *pthread\_cond\_timedwait()* or  
 32601 *pthread\_cond\_wait()* operations on the same condition variable is undefined; that is, a condition  
 32602 variable becomes bound to a unique mutex when a thread waits on the condition variable, and  
 32603 this (dynamic) binding ends when the wait returns.

32604 A condition wait (whether timed or not) is a cancelation point. When the cancelability enable  
 32605 state of a thread is set to `PTHREAD_CANCEL_DEFERRED`, a side effect of acting upon a  
 32606 cancelation request while in a condition wait is that the mutex is (in effect) re-acquired before  
 32607 calling the first cancelation cleanup handler. The effect is as if the thread were unblocked,  
 32608 allowed to execute up to the point of returning from the call to *pthread\_cond\_timedwait()* or  
 32609 *pthread\_cond\_wait()*, but at that point notices the cancelation request and instead of returning to  
 32610 the caller of *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()*, starts the thread cancelation  
 32611 activities, which includes calling cancelation cleanup handlers.

32612 A thread that has been unblocked because it has been canceled while blocked in a call to  
 32613 *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()* shall not consume any condition signal that  
 32614 may be directed concurrently at the condition variable if there are other threads blocked on the  
 32615 condition variable.

32616 The *pthread\_cond\_timedwait()* function is the same as *pthread\_cond\_wait()* except that an error is  
 32617 returned if the absolute time specified by *abstime* passes (that is, system time equals or exceeds  
 32618 *abstime*) before the condition *cond* is signaled or broadcasted, or if the absolute time specified by  
 32619 *abstime* has already been passed at the time of the call. If the Clock Selection option is supported,  
 32620 the condition variable shall have a clock attribute which specifies the clock that shall be used to

32621 measure the time specified by the *abstime* argument. When such timeouts occur,  
32622 *pthread\_cond\_timedwait()* shall nonetheless release and re-acquire the mutex referenced by *mutex*.  
32623 The *pthread\_cond\_timedwait()* function is also a cancellation point.

32624 If a signal is delivered to a thread waiting for a condition variable, upon return from the signal  
32625 handler the thread resumes waiting for the condition variable as if it was not interrupted, or it  
32626 shall return zero due to spurious wakeup.

#### 32627 RETURN VALUE

32628 Except in the case of [ETIMEDOUT], all these error checks shall act as if they were performed  
32629 immediately at the beginning of processing for the function and shall cause an error return, in  
32630 effect, prior to modifying the state of the mutex specified by *mutex* or the condition variable  
32631 specified by *cond*.

32632 Upon successful completion, a value of zero shall be returned; otherwise, an error number shall  
32633 be returned to indicate the error.

#### 32634 ERRORS

32635 The *pthread\_cond\_timedwait()* function shall fail if:

32636 [ETIMEDOUT] The time specified by *abstime* to *pthread\_cond\_timedwait()* has passed.

32637 The *pthread\_cond\_timedwait()* and *pthread\_cond\_wait()* functions may fail if:

32638 [EINVAL] The value specified by *cond*, *mutex*, or *abstime* is invalid.

32639 [EINVAL] Different mutexes were supplied for concurrent *pthread\_cond\_timedwait()* or  
32640 *pthread\_cond\_wait()* operations on the same condition variable.

32641 [EPERM] The mutex was not owned by the current thread at the time of the call.

32642 These functions shall not return an error code of [EINTR].

#### 32643 EXAMPLES

32644 None.

#### 32645 APPLICATION USAGE

32646 None.

#### 32647 RATIONALE

##### 32648 Condition Wait Semantics

32649 It is important to note that when *pthread\_cond\_wait()* and *pthread\_cond\_timedwait()* return  
32650 without error, the associated predicate may still be false. Similarly, when  
32651 *pthread\_cond\_timedwait()* returns with the timeout error, the associated predicate may be true  
32652 due to an unavoidable race between the expiration of the timeout and the predicate state change.

32653 Some implementations, particularly on a multiprocessor, may sometimes cause multiple threads  
32654 to wake up when the condition variable is signaled simultaneously on different processors.

32655 In general, whenever a condition wait returns, the thread has to re-evaluate the predicate  
32656 associated with the condition wait to determine whether it can safely proceed, should wait  
32657 again, or should declare a timeout. A return from the wait does not imply that the associated  
32658 predicate is either true or false.

32659 It is thus recommended that a condition wait be enclosed in the equivalent of a “while loop”  
32660 that checks the predicate.

32661 **Timed Wait Semantics**

32662 An absolute time measure was chosen for specifying the timeout parameter for two reasons.  
 32663 First, a relative time measure can be easily implemented on top of a function that specifies  
 32664 absolute time, but there is a race condition associated with specifying an absolute timeout on top  
 32665 of a function that specifies relative timeouts. For example, assume that `clock_gettime()` returns  
 32666 the current time and `cond_relative_timed_wait()` uses relative timeouts:

```
32667 clock_gettime(CLOCK_REALTIME, &now)
32668 reltime = sleep_til_this_absolute_time -now;
32669 cond_relative_timed_wait(c, m, &reltime);
```

32670 If the thread is preempted between the first statement and the last statement, the thread blocks  
 32671 for too long. Blocking, however, is irrelevant if an absolute timeout is used. An absolute timeout  
 32672 also need not be recomputed if it is used multiple times in a loop, such as that enclosing a  
 32673 condition wait.

32674 For cases when the system clock is advanced discontinuously by an operator, it is expected that  
 32675 implementations process any timed wait expiring at an intervening time as if that time had  
 32676 actually occurred.

32677 **Cancelation and Condition Wait**

32678 A condition wait, whether timed or not, is a cancelation point. That is, the functions  
 32679 `pthread_cond_wait()` or `pthread_cond_timedwait()` are points where a pending (or concurrent)  
 32680 cancelation request is noticed. The reason for this is that an indefinite wait is possible at these  
 32681 points—whatever event is being waited for, even if the program is totally correct, might never  
 32682 occur; for example, some input data being awaited might never be sent. By making condition  
 32683 wait a cancelation point, the thread can be canceled and perform its cancelation cleanup handler  
 32684 even though it may be stuck in some indefinite wait.

32685 A side effect of acting on a cancelation request while a thread is blocked on a condition variable  
 32686 is to re-acquire the mutex before calling any of the cancelation cleanup handlers. This is done in  
 32687 order to ensure that the cancelation cleanup handler is executed in the same state as the critical  
 32688 code that lies both before and after the call to the condition wait function. This rule is also  
 32689 required when interfacing to POSIX threads from languages, such as Ada or C++, which may  
 32690 choose to map cancelation onto a language exception; this rule ensures that each exception  
 32691 handler guarding a critical section can always safely depend upon the fact that the associated  
 32692 mutex has already been locked regardless of exactly where within the critical section the  
 32693 exception was raised. Without this rule, there would not be a uniform rule that exception  
 32694 handlers could follow regarding the lock, and so coding would become very cumbersome.

32695 Therefore, since *some* statement has to be made regarding the state of the lock when a  
 32696 cancelation is delivered during a wait, a definition has been chosen that makes application  
 32697 coding most convenient and error free.

32698 When acting on a cancelation request while a thread is blocked on a condition variable, the  
 32699 implementation is required to ensure that the thread does not consume any condition signals  
 32700 directed at that condition variable if there are any other threads waiting on that condition  
 32701 variable. This rule is specified in order to avoid deadlock conditions that could occur if these two  
 32702 independent requests (one acting on a thread and the other acting on the condition variable)  
 32703 were not processed independently.

**32704 Performance of Mutexes and Condition Variables**

32705 Mutexes are expected to be locked only for a few instructions. This practice is almost  
32706 automatically enforced by the desire of programmers to avoid long serial regions of execution  
32707 (which would reduce total effective parallelism).

32708 When using mutexes and condition variables, one tries to ensure that the usual case is to lock the  
32709 mutex, access shared data, and unlock the mutex. Waiting on a condition variable should be a  
32710 relatively rare situation. For example, when implementing a read-write lock, code that acquires a  
32711 read-lock typically needs only to increment the count of readers (under mutual-exclusion) and  
32712 return. The calling thread would actually wait on the condition variable only when there is  
32713 already an active writer. So the efficiency of a synchronization operation is bounded by the cost  
32714 of mutex lock/unlock and not by condition wait. Note that in the usual case there is no context  
32715 switch.

32716 This is not to say that the efficiency of condition waiting is unimportant. Since there needs to be  
32717 at least one context switch per Ada rendezvous, the efficiency of waiting on a condition variable  
32718 is important. The cost of waiting on a condition variable should be little more than the minimal  
32719 cost for a context switch plus the time to unlock and lock the mutex.

**32720 Features of Mutexes and Condition Variables**

32721 It had been suggested that the mutex acquisition and release be decoupled from condition wait.  
32722 This was rejected because it is the combined nature of the operation that, in fact, facilitates  
32723 realtime implementations. Those implementations can atomically move a high-priority thread  
32724 between the condition variable and the mutex in a manner that is transparent to the caller. This  
32725 can prevent extra context switches and provide more deterministic acquisition of a mutex when  
32726 the waiting thread is signaled. Thus, fairness and priority issues can be dealt with directly by the  
32727 scheduling discipline. Furthermore, the current condition wait operation matches existing  
32728 practice.

**32729 Scheduling Behavior of Mutexes and Condition Variables**

32730 Synchronization primitives that attempt to interfere with scheduling policy by specifying an  
32731 ordering rule are considered undesirable. Threads waiting on mutexes and condition variables  
32732 are selected to proceed in an order dependent upon the scheduling policy rather than in some  
32733 fixed order (for example, FIFO or priority). Thus, the scheduling policy determines which  
32734 thread(s) are awakened and allowed to proceed.

**32735 Timed Condition Wait**

32736 The *pthread\_cond\_timedwait()* function allows an application to give up waiting for a particular  
32737 condition after a given amount of time. An example of its use follows:

```
32738 (void) pthread_mutex_lock(&t.mn);
32739 t.waiters++;
32740 clock_gettime(CLOCK_REALTIME, &ts);
32741 ts.tv_sec += 5;
32742 rc = 0;
32743 while (! mypredicate(&t) && rc == 0)
32744 rc = pthread_cond_timedwait(&t.cond, &t.mn, &ts);
32745 t.waiters--;
32746 if (rc == 0) setmystate(&t);
32747 (void) pthread_mutex_unlock(&t.mn);
```

32748 By making the timeout parameter absolute, it does not need to be recomputed each time the  
32749 program checks its blocking predicate. If the timeout was relative, it would have to be  
32750 recomputed before each call. This would be especially difficult since such code would need to  
32751 take into account the possibility of extra wakeups that result from extra broadcasts or signals on  
32752 the condition variable that occur before either the predicate is true or the timeout is due.

32753 **FUTURE DIRECTIONS**

32754 None.

32755 **SEE ALSO**

32756 *pthread\_cond\_signal()*, *pthread\_cond\_broadcast()*, the Base Definitions volume of  
32757 IEEE Std. 1003.1-200x, <pthread.h>

32758 **CHANGE HISTORY**

32759 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32760 **Issue 6**

32761 The *pthread\_cond\_timedwait()* and *pthread\_cond\_wait()* functions are marked as part of the  
32762 Threads option.

32763 The Open Group corrigenda item U021/9 has been applied, correcting the prototype for the  
32764 *pthread\_cond\_wait()* function.

32765 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by adding semantics  
32766 for the Clock Selection option.

32767 The ERRORS section has an additional case for [EPERM] in response to IEEE PASC  
32768 Interpretation 1003.1c #28.

32769 The **restrict** keyword is added to the *pthread\_cond\_timedwait()* and *pthread\_cond\_wait()*  
32770 prototypes for alignment with the ISO/IEC 9899:1999 standard.

32771 **NAME**

32772 pthread\_cond\_wait — wait on a condition

32773 **SYNOPSIS**

32774 THR #include <pthread.h>

32775 int pthread\_cond\_wait(pthread\_cond\_t \*restrict cond,  
32776 pthread\_mutex\_t \*restrict mutex);

32777

32778 **DESCRIPTION**

32779 Refer to *pthread\_cond\_timedwait()*.

32780 **NAME**

32781 pthread\_condattr\_destroy, pthread\_condattr\_init — destroy and initialize condition variable  
 32782 attributes object

32783 **SYNOPSIS**

```
32784 THR #include <pthread.h>
```

```
32785 int pthread_condattr_destroy(pthread_condattr_t *attr);
```

```
32786 int pthread_condattr_init(pthread_condattr_t *attr);
```

```
32787
```

32788 **DESCRIPTION**

32789 The *pthread\_condattr\_destroy()* function destroys a condition variable attributes object; the object  
 32790 becomes, in effect, uninitialized. An implementation may cause *pthread\_condattr\_destroy()* to set  
 32791 the object referenced by *attr* to an invalid value. A destroyed condition variable attributes object  
 32792 can be re-initialized using *pthread\_condattr\_init()*; the results of otherwise referencing the object  
 32793 after it has been destroyed are undefined.

32794 The *pthread\_condattr\_init()* function initializes a condition variable attributes object *attr* with the  
 32795 default value for all of the attributes defined by the implementation.

32796 Attempting to initialize an already initialized condition variable attributes object results in  
 32797 undefined behavior.

32798 After a condition variable attributes object has been used to initialize one or more condition  
 32799 variables, any function affecting the attributes object (including destruction) does not affect any  
 32800 previously initialized condition variables.

32801 Additional attributes, their default values, and the names of the associated functions to get and  
 32802 set those attribute values are implementation-defined.

32803 **RETURN VALUE**

32804 If successful, the *pthread\_condattr\_destroy()* and *pthread\_condattr\_init()* functions shall return  
 32805 zero; otherwise, an error number shall be returned to indicate the error.

32806 **ERRORS**

32807 The *pthread\_condattr\_destroy()* function may fail if:

32808 [EINVAL] The value specified by *attr* is invalid.

32809 The *pthread\_condattr\_init()* function shall fail if:

32810 [ENOMEM] Insufficient memory exists to initialize the condition variable attributes object.

32811 These functions shall not return an error code of [EINTR].

32812 **EXAMPLES**

32813 None.

32814 **APPLICATION USAGE**

32815 None.

32816 **RATIONALE**

32817 See *pthread\_attr\_init()* and *pthread\_mutex\_init()*.

32818 A *process-shared* attribute has been defined for condition variables for the same reason it has been  
 32819 defined for mutexes.

32820 **FUTURE DIRECTIONS**

32821 None.

32822 **SEE ALSO**

32823 *pthread\_cond\_destroy()*, *pthread\_condattr\_getpshared()*, *pthread\_create()*, *pthread\_mutex\_destroy()*,  
32824 the Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

32825 **CHANGE HISTORY**

32826 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32827 **Issue 6**

32828 The *pthread\_condattr\_destroy()* and *pthread\_condattr\_init()* functions are marked as part of the  
32829 Threads option.

32830 **NAME**

32831 pthread\_condattr\_getclock, pthread\_condattr\_setclock — get and set the clock selection  
 32832 condition variable attribute

32833 **SYNOPSIS**

32834 THR CS #include <pthread.h>

```
32835 int pthread_condattr_getclock(const pthread_condattr_t *restrict attr,
32836 clockid_t *restrict clock_id);
32837 int pthread_condattr_setclock(pthread_condattr_t *attr,
32838 clockid_t clock_id);
32839
```

32840 **DESCRIPTION**

32841 The *pthread\_condattr\_getclock()* function obtains the value of the clock attribute from the  
 32842 attributes object referenced by *attr*. The *pthread\_condattr\_setclock()* function is used to set the  
 32843 clock attribute in an initialized attributes object referenced by *attr*. If *pthread\_condattr\_setclock()*  
 32844 is called with a *clock\_id* argument that refers to a CPU-time clock, the call shall fail.

32845 The clock attribute is the clock ID of the clock that shall be used to measure the timeout service  
 32846 of *pthread\_cond\_timedwait()*. The default value of the clock attribute shall refer to the system  
 32847 clock.

32848 **RETURN VALUE**

32849 If successful, the *pthread\_condattr\_getclock()* function shall return zero and store the value of the  
 32850 clock attribute of *attr* into the object referenced by the *clock\_id* argument. Otherwise, an error  
 32851 number shall be returned to indicate the error.

32852 If successful, the *pthread\_condattr\_setclock()* function shall return zero; otherwise, an error  
 32853 number shall be returned to indicate the error.

32854 **ERRORS**

32855 These functions may fail if:

32856 [EINVAL] The value specified by *attr* is invalid.

32857 The *pthread\_condattr\_setclock()* function may fail if:

32858 [EINVAL] The value specified by *clock\_id* does not refer to a known clock, or is a CPU-  
 32859 time clock.

32860 These functions shall not return an error code of [EINTR].

32861 **EXAMPLES**

32862 None.

32863 **APPLICATION USAGE**

32864 None.

32865 **RATIONALE**

32866 None.

32867 **FUTURE DIRECTIONS**

32868 None.

32869 **SEE ALSO**

32870 *pthread\_cond\_init()*, *pthread\_cond\_timedwait()*, *pthread\_condattr\_destroy()*,  
 32871 *pthread\_condattr\_getshared()* (on page 1553), *pthread\_condattr\_init()*,  
 32872 *pthread\_condattr\_setshared()* (on page 1557), *pthread\_create()*, *pthread\_mutex\_init()*, the Base  
 32873 Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

32874 **CHANGE HISTORY**

32875 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

32876 **NAME**

32877 pthread\_condattr\_getpshared, pthread\_condattr\_setpshared — get and set the process-shared  
 32878 condition variable attributes

32879 **SYNOPSIS**

```
32880 THR TSH #include <pthread.h>
32881
32882 int pthread_condattr_getpshared(const pthread_condattr_t *restrict attr,
32883 int *restrict pshared);
32884 int pthread_condattr_setpshared(pthread_condattr_t *attr,
32885 int pshared);
```

32886 **DESCRIPTION**

32887 The *pthread\_condattr\_getpshared()* function obtains the value of the *process-shared* attribute from  
 32888 the attributes object referenced by *attr*. The *pthread\_condattr\_setpshared()* function is used to set  
 32889 the *process-shared* attribute in an initialized attributes object referenced by *attr*.

32890 The *process-shared* attribute is set to `PTHREAD_PROCESS_SHARED` to permit a condition  
 32891 variable to be operated upon by any thread that has access to the memory where the condition  
 32892 variable is allocated, even if the condition variable is allocated in memory that is shared by  
 32893 multiple processes. If the *process-shared* attribute is `PTHREAD_PROCESS_PRIVATE`, the  
 32894 condition variable is only operated upon by threads created within the same process as the  
 32895 thread that initialized the condition variable; if threads of differing processes attempt to operate  
 32896 on such a condition variable, the behavior is undefined. The default value of the attribute is  
 32897 `PTHREAD_PROCESS_PRIVATE`.

32898 Additional attributes, their default values, and the names of the associated functions to get and  
 32899 set those attribute values are implementation-defined.

32900 **RETURN VALUE**

32901 If successful, the *pthread\_condattr\_setpshared()* function shall return zero; otherwise, an error  
 32902 number shall be returned to indicate the error.

32903 If successful, the *pthread\_condattr\_getpshared()* function shall return zero and store the value of  
 32904 the *process-shared* attribute of *attr* into the object referenced by the *pshared* parameter. Otherwise,  
 32905 an error number shall be returned to indicate the error.

32906 **ERRORS**

32907 The *pthread\_condattr\_getpshared()* and *pthread\_condattr\_setpshared()* functions may fail if:

32908 [EINVAL] The value specified by *attr* is invalid.

32909 The *pthread\_condattr\_setpshared()* function may fail if:

32910 [EINVAL] The new value specified for the attribute is outside the range of legal values  
 32911 for that attribute.

32912 These functions shall not return an error code of [EINTR].

32913 **EXAMPLES**

32914 None.

32915 **APPLICATION USAGE**

32916 None.

32917 **RATIONALE**

32918 None.

32919 **FUTURE DIRECTIONS**

32920 None.

32921 **SEE ALSO**

32922 *pthread\_create()*, *pthread\_cond\_destroy()*, *pthread\_condattr\_destroy()*, *pthread\_mutex\_destroy()*, the  
32923 Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

32924 **CHANGE HISTORY**

32925 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

32926 **Issue 6**

32927 The *pthread\_condattr\_getpshared()* and *pthread\_condattr\_setpshared()* functions are marked as part  
32928 of the Threads and Thread Process-Shared Synchronization options.

32929 The **restrict** keyword is added to the *pthread\_condattr\_getpshared()* prototype for alignment with  
32930 the ISO/IEC 9899:1999 standard.

32931 **NAME**

32932 pthread\_condattr\_init — initialize condition variable attributes object

32933 **SYNOPSIS**

32934 THR #include &lt;pthread.h&gt;

32935 int pthread\_condattr\_init(pthread\_condattr\_t \*attr);

32936

32937 **DESCRIPTION**32938 Refer to *pthread\_condattr\_destroy()*.

32939 **NAME**

32940 pthread\_condattr\_setclock — set the clock selection condition variable attribute

32941 **SYNOPSIS**

32942 THR CS #include <pthread.h>

32943 int pthread\_condattr\_setclock(pthread\_condattr\_t \*attr,  
32944 clockid\_t clock\_id);

32945

32946 **DESCRIPTION**

32947 Refer to *pthread\_condattr\_getclock()*.

32948 **NAME**

32949 pthread\_condattr\_setpshared — set the process-shared condition variable attributes

32950 **SYNOPSIS**

```
32951 THR TSH #include <pthread.h>
```

```
32952 int pthread_condattr_setpshared(pthread_condattr_t *attr,
32953 int pshared);
```

32954

32955 **DESCRIPTION**

32956 Refer to *pthread\_condattr\_getpshared()*.

## 32957 NAME

32958 pthread\_create — thread creation

## 32959 SYNOPSIS

32960 THR #include &lt;pthread.h&gt;

```
32961 int pthread_create(pthread_t *restrict thread,
32962 const pthread_attr_t *restrict attr,
32963 void *(*start_routine)(void*), void *arg);
32964
```

## 32965 DESCRIPTION

32966 The *pthread\_create()* function is used to create a new thread, with attributes specified by *attr*,  
 32967 within a process. If *attr* is NULL, the default attributes are used. If the attributes specified by *attr*  
 32968 are modified later, the thread's attributes are not affected. Upon successful completion,  
 32969 *pthread\_create()* shall store the ID of the created thread in the location referenced by *thread*.

32970 The thread is created executing *start\_routine* with *arg* as its sole argument. If the *start\_routine*  
 32971 returns, the effect shall be as if there was an implicit call to *pthread\_exit()* using the return value  
 32972 of *start\_routine* as the exit status. Note that the thread in which *main()* was originally invoked  
 32973 differs from this. When it returns from *main()*, the effect shall be as if there was an implicit call  
 32974 to *exit()* using the return value of *main()* as the exit status.

32975 The signal state of the new thread shall be initialized as follows:

- 32976 • The signal mask shall be inherited from the creating thread.
- 32977 • The set of signals pending for the new thread shall be empty.

32978 If *pthread\_create()* fails, no new thread is created and the contents of the location referenced by  
 32979 *thread* are undefined.

32980 TCT If `_POSIX_THREAD_CPUTIME` is defined, the new thread shall have a CPU-time clock  
 32981 accessible, and the initial value of this clock shall be set to zero.

## 32982 RETURN VALUE

32983 If successful, the *pthread\_create()* function shall return zero; otherwise, an error number shall be  
 32984 returned to indicate the error.

## 32985 ERRORS

32986 The *pthread\_create()* function shall fail if:

32987 [EAGAIN] The system lacked the necessary resources to create another thread, or the  
 32988 system-imposed limit on the total number of threads in a process  
 32989 PTHREAD\_THREADS\_MAX would be exceeded.

32990 [EINVAL] The value specified by *attr* is invalid.

32991 [EPERM] The caller does not have appropriate permission to set the required  
 32992 scheduling parameters or scheduling policy.

32993 The *pthread\_create()* function shall not return an error code of [EINTR].

32994 **EXAMPLES**

32995 None.

32996 **APPLICATION USAGE**

32997 None.

32998 **RATIONALE**

32999 A suggested alternative to *pthread\_create()* would be to define two separate operations: create  
 33000 and start. Some applications would find such behavior more natural. Ada, in particular,  
 33001 separates the “creation” of a task from its “activation”.

33002 Splitting the operation was rejected by the standard developers for many reasons:

- 33003 • The number of calls required to start a thread would increase from one to two and thus place  
 33004 an additional burden on applications that do not require the additional synchronization. The  
 33005 second call, however, could be avoided by the additional complication of a start-up state  
 33006 attribute.
- 33007 • An extra state would be introduced: “created but not started”. This would require the  
 33008 standard to specify the behavior of the thread operations when the target has not yet started  
 33009 executing.
- 33010 • For those applications that require such behavior, it is possible to simulate the two separate  
 33011 steps with the facilities that are currently provided. The *start\_routine()* can synchronize by  
 33012 waiting on a condition variable that is signaled by the start operation.

33013 An Ada implementor can choose to create the thread at either of two points in the Ada program:  
 33014 when the task object is created, or when the task is activated (generally at a “begin”). If the first  
 33015 approach is adopted, the *start\_routine()* needs to wait on a condition variable to receive the  
 33016 order to begin “activation”. The second approach requires no such condition variable or extra  
 33017 synchronization. In either approach, a separate Ada task control block would need to be created  
 33018 when the task object is created to hold rendezvous queues, and so on.

33019 An extension of the preceding model would be to allow the state of the thread to be modified  
 33020 between the create and start. This would allow the thread attributes object to be eliminated. This  
 33021 has been rejected because:

- 33022 • All state in the thread attributes object has to be able to be set for the thread. This would  
 33023 require the definition of functions to modify thread attributes. There would be no reduction  
 33024 in the number of function calls required to set up the thread. In fact, for an application that  
 33025 creates all threads using identical attributes, the number of function calls required to set up  
 33026 the threads would be dramatically increased. Use of a thread attributes object permits the  
 33027 application to make one set of attribute setting function calls. Otherwise, the set of attribute  
 33028 setting function calls needs to be made for each thread creation.
- 33029 • Depending on the implementation architecture, functions to set thread state would require  
 33030 kernel calls, or for other implementation reasons would not be able to be implemented as  
 33031 macros, thereby increasing the cost of thread creation.
- 33032 • The ability for applications to segregate threads by class would be lost.

33033 Another suggested alternative uses a model similar to that for process creation, such as “thread  
 33034 fork”. The fork semantics would provide more flexibility and the “create” function can be  
 33035 implemented simply by doing a thread fork followed immediately by a call to the desired “start  
 33036 routine” for the thread. This alternative has these problems:

- 33037 • For many implementations, the entire stack of the calling thread would need to be  
 33038 duplicated, since in many architectures there is no way to determine the size of the calling  
 33039 frame.

33040           • Efficiency is reduced since at least some part of the stack has to be copied, even though in  
33041           most cases the thread never needs the copied context, since it merely calls the desired start  
33042           routine.

## 33043 FUTURE DIRECTIONS

33044           None.

## 33045 SEE ALSO

33046           *fork()*, *pthread\_exit()*, *pthread\_join()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
33047           <pthread.h>

## 33048 CHANGE HISTORY

33049           First released in Issue 5. Included for alignment with the POSIX Threads Extension.

## 33050 Issue 6

33051           The *pthread\_create()* function is marked as part of the Threads option.

33052           The following new requirements on POSIX implementations derive from alignment with the  
33053           Single UNIX Specification:

33054           • The [EPERM] mandatory error condition is added.

33055           The thread CPU-time clock semantics are added for alignment with IEEE Std. 1003.1d-1999.

33056           The **restrict** keyword is added to the *pthread\_create()* prototype for alignment with the  
33057           ISO/IEC 9899:1999 standard.

33058 **NAME**

33059 pthread\_detach — detach a thread

33060 **SYNOPSIS**

33061 THR #include &lt;pthread.h&gt;

33062 int pthread\_detach(pthread\_t thread);

33063

33064 **DESCRIPTION**

33065 The *pthread\_detach()* function is used to indicate to the implementation that storage for the  
 33066 thread *thread* can be reclaimed when that thread terminates. If *thread* has not terminated,  
 33067 *pthread\_detach()* shall not cause it to terminate. The effect of multiple *pthread\_detach()* calls on  
 33068 the same target thread is unspecified.

33069 **RETURN VALUE**

33070 If the call succeeds, *pthread\_detach()* shall return 0; otherwise, an error number shall be returned  
 33071 to indicate the error.

33072 **ERRORS**33073 The *pthread\_detach()* function shall fail if:

33074 [EINVAL] The implementation has detected that the value specified by *thread* does not  
 33075 refer to a joinable thread.

33076 [ESRCH] No thread could be found corresponding to that specified by the given thread  
 33077 ID.

33078 The *pthread\_detach()* function shall not return an error code of [EINTR].33079 **EXAMPLES**

33080 None.

33081 **APPLICATION USAGE**

33082 None.

33083 **RATIONALE**

33084 The *pthread\_join()* or *pthread\_detach()* functions should eventually be called for every thread that  
 33085 is created so that storage associated with the thread may be reclaimed.

33086 It has been suggested that a “detach” function is not necessary; the *detachstate* thread creation  
 33087 attribute is sufficient, since a thread need never be dynamically detached. However, need arises  
 33088 in at least two cases:

33089 1. In a cancellation handler for a *pthread\_join()* it is nearly essential to have a *pthread\_detach()*  
 33090 function in order to detach the thread on which *pthread\_join()* was waiting. Without it, it  
 33091 would be necessary to have the handler do another *pthread\_join()* to attempt to detach the  
 33092 thread, which would both delay the cancellation processing for an unbounded period and  
 33093 introduce a new call to *pthread\_join()*, which might itself need a cancellation handler. A  
 33094 dynamic detach is nearly essential in this case.

33095 2. In order to detach the “initial thread” (as may be desirable in processes that set up server  
 33096 threads).

33097 **FUTURE DIRECTIONS**

33098 None.

33099 **SEE ALSO**

33100 *pthread\_join()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

33101 **CHANGE HISTORY**

33102 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33103 **Issue 6**

33104 The *pthread\_detach()* function is marked as part of the Threads option.

33105 **NAME**

33106 pthread\_equal — compare thread IDs

33107 **SYNOPSIS**

33108 THR #include &lt;pthread.h&gt;

33109 int pthread\_equal(pthread\_t t1, pthread\_t t2);

33110

33111 **DESCRIPTION**33112 This function compares the thread IDs *t1* and *t2*.33113 **RETURN VALUE**33114 The *pthread\_equal()* function shall return a non-zero value if *t1* and *t2* are equal; otherwise, zero  
33115 shall be returned.33116 If either *t1* or *t2* are not valid thread IDs, the behavior is undefined.33117 **ERRORS**

33118 No errors are defined.

33119 The *pthread\_equal()* function shall not return an error code of [EINTR].33120 **EXAMPLES**

33121 None.

33122 **APPLICATION USAGE**

33123 None.

33124 **RATIONALE**33125 Implementations may choose to define a thread ID as a structure. This allows additional  
33126 flexibility and robustness over using an **int**. For example, a thread ID could include a sequence  
33127 number that allows detection of “dangling IDs” (copies of a thread ID that has been detached).  
33128 Because the C language does not support comparison on structure types, the *pthread\_equal()*  
33129 function is provided to compare thread IDs.33130 **FUTURE DIRECTIONS**

33131 None.

33132 **SEE ALSO**33133 *pthread\_create()*, *pthread\_self()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
33134 <pthread.h>33135 **CHANGE HISTORY**

33136 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33137 **Issue 6**33138 The *pthread\_equal()* function is marked as part of the Threads option.

33139 **NAME**

33140 pthread\_exit — thread termination

33141 **SYNOPSIS**

33142 THR #include &lt;pthread.h&gt;

33143 void pthread\_exit(void \*value\_ptr);

33144

33145 **DESCRIPTION**

33146 The *pthread\_exit()* function shall terminate the calling thread and make the value *value\_ptr*  
33147 available to any successful join with the terminating thread. Any cancellation cleanup handlers  
33148 that have been pushed and not yet popped shall be popped in the reverse order that they were  
33149 pushed and then executed. After all cancellation cleanup handlers have been executed, if the  
33150 thread has any thread-specific data, appropriate destructor functions shall be called in an  
33151 unspecified order. Thread termination does not release any application visible process resources,  
33152 including, but not limited to, mutexes and file descriptors, nor does it perform any process-level  
33153 cleanup actions, including, but not limited to, calling any *atexit()* routines that may exist.

33154 An implicit call to *pthread\_exit()* is made when a thread other than the thread in which *main()*  
33155 was first invoked returns from the start routine that was used to create it. The function's return  
33156 value serves as the thread's exit status.

33157 The behavior of *pthread\_exit()* is undefined if called from a cancellation cleanup handler or  
33158 destructor function that was invoked as a result of either an implicit or explicit call to  
33159 *pthread\_exit()*.

33160 After a thread has terminated, the result of access to local (auto) variables of the thread is  
33161 undefined. Thus, references to local variables of the exiting thread should not be used for the  
33162 *pthread\_exit()* *value\_ptr* parameter value.

33163 The process shall exit with an exit status of 0 after the last thread has been terminated. The  
33164 behavior shall be as if the implementation called *exit()* with a zero argument at thread  
33165 termination time.

33166 **RETURN VALUE**33167 The *pthread\_exit()* function cannot return to its caller.33168 **ERRORS**

33169 No errors are defined.

33170 The *pthread\_exit()* function shall not return an error code of [EINTR].33171 **EXAMPLES**

33172 None.

33173 **APPLICATION USAGE**

33174 None.

33175 **RATIONALE**

33176 The normal mechanism by which a thread terminates is to return from the routine that was  
33177 specified in the *pthread\_create()* call that started it. The *pthread\_exit()* function provides the  
33178 capability for a thread to terminate without requiring a return from the start routine of that  
33179 thread, thereby providing a function analogous to *exit()*.

33180 Regardless of the method of thread termination, any cancellation cleanup handlers that have  
33181 been pushed and not yet popped are executed, and the destructors for any existing thread-  
33182 specific data are executed. This volume of IEEE Std. 1003.1-200x requires that cancellation  
33183 cleanup handlers be popped and called in order. After all cancellation cleanup handlers have

33184 been executed, thread-specific data destructors are called, in an unspecified order, for each item  
33185 of thread-specific data that exists in the thread. This ordering is necessary because cancellation  
33186 cleanup handlers may rely on thread-specific data.

33187 As the meaning of the status is determined by the application (except when the thread has been  
33188 canceled, in which case it is PTHREAD\_CANCELED), the implementation has no idea what an  
33189 illegal status value is, which is why no address error checking is done.

33190 **FUTURE DIRECTIONS**

33191 None.

33192 **SEE ALSO**

33193 *exit()*, *pthread\_create()*, *pthread\_join()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
33194 <pthread.h>

33195 **CHANGE HISTORY**

33196 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33197 **Issue 6**

33198 The *pthread\_exit()* function is marked as part of the Threads option. |

33199 **NAME**

33200 pthread\_getconcurrency, pthread\_setconcurrency — get or set level of concurrency

33201 **SYNOPSIS**33202 XSI 

```
#include <pthread.h>
```

33203 

```
int pthread_getconcurrency(void);
```

33204 

```
int pthread_setconcurrency(int new_level);
```

33205

33206 **DESCRIPTION**

33207 Unbound threads in a process may or may not be required to be simultaneously active. By  
 33208 default, the threads implementation ensures that a sufficient number of threads are active so that  
 33209 the process can continue to make progress. While this conserves system resources, it may not  
 33210 produce the most effective level of concurrency.

33211 The *pthread\_setconcurrency()* function allows an application to inform the threads  
 33212 implementation of its desired concurrency level, *new\_level*. The actual level of concurrency  
 33213 provided by the implementation as a result of this function call is unspecified.

33214 If *new\_level* is zero, it causes the implementation to maintain the concurrency level at its  
 33215 discretion as if *pthread\_setconcurrency()* had never been called.

33216 The *pthread\_getconcurrency()* function shall return the value set by a previous call to the  
 33217 *pthread\_setconcurrency()* function. If the *pthread\_setconcurrency()* function was not previously  
 33218 called, this function shall return zero to indicate that the implementation is maintaining the  
 33219 concurrency level.

33220 When an application calls *pthread\_setconcurrency()* it is informing the implementation of its  
 33221 desired concurrency level. The implementation uses this as a hint, not a requirement.

33222 If an implementation does not support multiplexing of user threads on top of several kernel-  
 33223 scheduled entities, the *pthread\_setconcurrency()* and *pthread\_getconcurrency()* functions are  
 33224 provided for source code compatibility but they have no effect when called. To maintain the  
 33225 function semantics, the *new\_level* parameter is saved when *pthread\_setconcurrency()* is called so  
 33226 that a subsequent call to *pthread\_getconcurrency()* returns the same value.

33227 **RETURN VALUE**

33228 If successful, the *pthread\_setconcurrency()* function shall return zero; otherwise, an error number  
 33229 shall be returned to indicate the error.

33230 The *pthread\_getconcurrency()* function shall always return the concurrency level set by a previous  
 33231 call to *pthread\_setconcurrency()*. If the *pthread\_setconcurrency()* function has never been called,  
 33232 *pthread\_getconcurrency()* shall return zero.

33233 **ERRORS**33234 The *pthread\_setconcurrency()* function shall fail if:33235 [EINVAL] The value specified by *new\_level* is negative.33236 [EAGAIN] The value specific by *new\_level* would cause a system resource to be exceeded.

33237 These functions shall not return an error code of [EINTR].

33238 **EXAMPLES**

33239 None.

33240 **APPLICATION USAGE**

33241 Use of these functions changes the state of the underlying concurrency upon which the  
33242 application depends. Library developers are advised to not use the *pthread\_getconcurrency()* and  
33243 *pthread\_setconcurrency()* functions since their use may conflict with an applications use of these  
33244 functions.

33245 **RATIONALE**

33246 None.

33247 **FUTURE DIRECTIONS**

33248 None.

33249 **SEE ALSO**33250 The Base Definitions volume of IEEE Std. 1003.1-200x, <**pthread.h**>33251 **CHANGE HISTORY**

33252 First released in Issue 5.

33253 **NAME**

33254 pthread\_getcpuclockid — access a thread CPU-time clock (**REALTIME**)

33255 **SYNOPSIS**

33256 TCT #include <pthread.h>

33257 #include <time.h>

33258 int pthread\_getcpuclockid(pthread\_t *thread\_id*, clockid\_t \**clock\_id*);

33259

33260 **DESCRIPTION**

33261 The *pthread\_getcpuclockid()* function shall return in *clock\_id* the clock ID of the CPU-time clock of  
33262 the thread specified by *thread\_id*, if the thread specified by *thread\_id* exists.

33263 **RETURN VALUE**

33264 Upon successful completion, *pthread\_getcpuclockid()* shall return zero; otherwise, an error  
33265 number shall be returned to indicate the error.

33266 **ERRORS**

33267 The *pthread\_getcpuclockid()* function may fail if:

33268 [ESRCH] The value specified by *thread\_id* does not refer to an existing thread.

33269 **EXAMPLES**

33270 None.

33271 **APPLICATION USAGE**

33272 The *pthread\_getcpuclockid()* function is part of the Thread CPU-Time Clocks option and need not  
33273 be provided on all implementations.

33274 **RATIONALE**

33275 None.

33276 **FUTURE DIRECTIONS**

33277 None.

33278 **SEE ALSO**

33279 *clock\_getcpuclockid()*, *clock\_getres()*, *timer\_create()*, the Base Definitions volume of  
33280 IEEE Std. 1003.1-200x, <pthread.h>, <time.h>

33281 **CHANGE HISTORY**

33282 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

33283 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

## 33284 NAME

33285 pthread\_getschedparam, pthread\_setschedparam — dynamic thread scheduling parameters  
 33286 access (**REALTIME THREADS**)

## 33287 SYNOPSIS

```
33288 TPS #include <pthread.h>
33289
33289 int pthread_getschedparam(pthread_t thread, int *restrict policy,
33290 struct sched_param *restrict param);
33291 int pthread_setschedparam(pthread_t thread, int policy,
33292 const struct sched_param *param);
33293
```

## 33294 DESCRIPTION

33295 The *pthread\_getschedparam()* and *pthread\_setschedparam()* functions allow the scheduling policy  
 33296 and scheduling parameters of individual threads within a multi-threaded process to be retrieved  
 33297 and set. For SCHED\_FIFO and SCHED\_RR, the only required member of the **sched\_param**  
 33298 structure is the priority *sched\_priority*. For SCHED\_OTHER, the affected scheduling parameters  
 33299 are implementation-defined.

33300 The *pthread\_getschedparam()* function shall retrieve the scheduling policy and scheduling  
 33301 parameters for the thread whose thread ID is given by *thread* and shall store those values in  
 33302 *policy* and *param*, respectively. The priority value returned from *pthread\_getschedparam()* shall be  
 33303 the value specified by the most recent *pthread\_setschedparam()* or *pthread\_create()* call affecting  
 33304 the target thread. It shall not reflect any temporary adjustments to its priority as a result of any  
 33305 priority inheritance or ceiling functions. The *pthread\_setschedparam()* function sets the scheduling  
 33306 policy and associated scheduling parameters for the thread whose thread ID is given by *thread* to  
 33307 the policy and associated parameters provided in *policy* and *param*, respectively.

33308 The *policy* parameter may have the value SCHED\_OTHER, SCHED\_FIFO, or SCHED\_RR. The  
 33309 scheduling parameters for the SCHED\_OTHER policy are implementation-defined. The  
 33310 SCHED\_FIFO and SCHED\_RR policies shall have a single scheduling parameter, *priority*.

33311 TSP If **\_POSIX\_THREAD\_SPORADIC\_SERVER** is defined, then the *policy* argument may have the  
 33312 value SCHED\_SPORADIC, with the exception for the *pthread\_setschedparam()* function that if the  
 33313 scheduling policy was not SCHED\_SPORADIC at the time of the call, it is implementation-  
 33314 defined whether the function is supported; in other words, the implementation need not allow  
 33315 the application to dynamically change the scheduling policy to SCHED\_SPORADIC. The  
 33316 sporadic server scheduling policy has the associated parameters *sched\_ss\_low\_priority*,  
 33317 *sched\_ss\_repl\_period*, *sched\_ss\_init\_budget*, *sched\_priority*, and *sched\_ss\_max\_repl*. The specified  
 33318 *sched\_ss\_repl\_period* shall be greater than or equal to the specified *sched\_ss\_init\_budget* for the  
 33319 function to succeed; if it is not, then the function shall fail. The value of *sched\_ss\_max\_repl* shall  
 33320 be within the inclusive range [1,{SS\_REPL\_MAX}] for the function to succeed; if not, the function  
 33321 shall fail.

33322 If the *pthread\_setschedparam()* function fails, no scheduling parameters are changed for the target  
 33323 thread.

## 33324 RETURN VALUE

33325 If successful, the *pthread\_getschedparam()* and *pthread\_setschedparam()* functions shall return zero;  
 33326 otherwise, an error number shall be returned to indicate the error.

## 33327 ERRORS

33328 The *pthread\_getschedparam()* function may fail if:

33329 [ESRCH] The value specified by *thread* does not refer to a existing thread.

- 33330 The *pthread\_setschedparam()* function may fail if:
- 33331 [EINVAL] The value specified by *policy* or one of the scheduling parameters associated  
33332 with the scheduling policy *policy* is invalid.
- 33333 [ENOTSUP] An attempt was made to set the policy or scheduling parameters to an  
33334 unsupported value.
- 33335 TSP [ENOTSUP] An attempt was made to dynamically change the scheduling policy to  
33336 SCHED\_SPORADIC, and the implementation does not support this change.
- 33337 [EPERM] The caller does not have the appropriate permission to set either the  
33338 scheduling parameters or the scheduling policy of the specified thread.
- 33339 [EPERM] The implementation does not allow the application to modify one of the  
33340 parameters to the value specified.
- 33341 [ESRCH] The value specified by *thread* does not refer to a existing thread.
- 33342 These functions shall not return an error code of [EINTR].
- 33343 **EXAMPLES**
- 33344 None.
- 33345 **APPLICATION USAGE**
- 33346 None.
- 33347 **RATIONALE**
- 33348 None.
- 33349 **FUTURE DIRECTIONS**
- 33350 None.
- 33351 **SEE ALSO**
- 33352 *sched\_getparam()*, *sched\_getscheduler()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
33353 <pthread.h>, <sched.h>
- 33354 **CHANGE HISTORY**
- 33355 First released in Issue 5. Included for alignment with the POSIX Threads Extension.
- 33356 **Issue 6**
- 33357 The *pthread\_getschedparam()* and *pthread\_setschedparam()* functions are marked as part of the  
33358 Thread Execution Scheduling option.
- 33359 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
33360 implementation does not support the Thread Execution Scheduling option.
- 33361 The Open Group corrigenda item U026/2 has been applied correcting the prototype for the  
33362 *pthread\_setschedparam()* function so that its second argument is of type **int**.
- 33363 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std. 1003.1d-1999.
- 33364 The **restrict** keyword is added to the *pthread\_getschedparam()* prototype for alignment with the  
33365 ISO/IEC 9899:1999 standard.
- 33366 The Open Group corrigenda item U047/1 has been applied.

33367 **NAME**

33368 pthread\_getspecific, pthread\_setspecific — thread-specific data management

33369 **SYNOPSIS**

33370 THR #include &lt;pthread.h&gt;

33371 void \*pthread\_getspecific(pthread\_key\_t key);

33372 int pthread\_setspecific(pthread\_key\_t key, const void \*value);

33373

33374 **DESCRIPTION**33375 The *pthread\_getspecific()* function shall return the value currently bound to the specified *key* on  
33376 behalf of the calling thread.33377 The *pthread\_setspecific()* function shall associate a thread-specific *value* with a *key* obtained via a  
33378 previous call to *pthread\_key\_create()*. Different threads may bind different values to the same  
33379 key. These values are typically pointers to blocks of dynamically allocated memory that have  
33380 been reserved for use by the calling thread.33381 The effect of calling *pthread\_getspecific()* or *pthread\_setspecific()* with a *key* value not obtained  
33382 from *pthread\_key\_create()* or after *key* has been deleted with *pthread\_key\_delete()* is undefined.33383 Both *pthread\_getspecific()* and *pthread\_setspecific()* may be called from a thread-specific data  
33384 destructor function. A call to *pthread\_getspecific()* for the thread-specific data key being  
33385 destroyed shall return the value NULL, unless the value is changed (after the destructor starts)  
33386 by a call to *pthread\_setspecific()*. Calling *pthread\_setspecific()* thread-specific data destructor  
33387 routine may result either in lost storage (after at least PTHREAD\_DESTRUCTOR\_ITERATIONS)  
33388 or an infinite loop.

33389 Both functions may be implemented as macros.

33390 **RETURN VALUE**33391 The *pthread\_getspecific()* function shall return the thread-specific data value associated with the  
33392 given *key*. If no thread-specific data value is associated with *key*, then the value NULL shall be  
33393 returned.33394 If successful, the *pthread\_setspecific()* function shall return zero; otherwise, an error number shall  
33395 be returned to indicate the error.33396 **ERRORS**33397 No errors are returned from *pthread\_getspecific()*.33398 The *pthread\_setspecific()* function shall fail if:

33399 [ENOMEM] Insufficient memory exists to associate the value with the key.

33400 The *pthread\_setspecific()* function may fail if:

33401 [EINVAL] The key value is invalid.

33402 These functions shall not return an error code of [EINTR].

33403 **EXAMPLES**

33404 None.

33405 **APPLICATION USAGE**

33406 None.

33407 **RATIONALE**

33408 Performance and ease-of-use of *pthread\_getspecific()* is critical for functions that rely on  
33409 maintaining state in thread-specific data. Since no errors are required to be detected by it, and  
33410 since the only error that could be detected is the use of an invalid key, the function to  
33411 *pthread\_getspecific()* has been designed to favor speed and simplicity over error reporting.

33412 **FUTURE DIRECTIONS**

33413 None.

33414 **SEE ALSO**

33415 *pthread\_key\_create()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**pthread.h**>

33416 **CHANGE HISTORY**

33417 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33418 **Issue 6**

33419 The *pthread\_getspecific()* and *pthread\_setspecific()* functions are marked as part of the Threads  
33420 option.

33421 **NAME**

33422 pthread\_join — wait for thread termination

33423 **SYNOPSIS**

33424 THR #include &lt;pthread.h&gt;

33425 int pthread\_join(pthread\_t thread, void \*\*value\_ptr);

33426

33427 **DESCRIPTION**

33428 The *pthread\_join()* function shall suspend execution of the calling thread until the target *thread*  
 33429 terminates, unless the target *thread* has already terminated. On return from a successful  
 33430 *pthread\_join()* call with a non-NULL *value\_ptr* argument, the value passed to *pthread\_exit()* by  
 33431 the terminating thread shall be made available in the location referenced by *value\_ptr*. When a  
 33432 *pthread\_join()* returns successfully, the target thread has been terminated. The results of multiple  
 33433 simultaneous calls to *pthread\_join()* specifying the same target thread are undefined. If the  
 33434 thread calling *pthread\_join()* is canceled, then the target thread shall not be detached.

33435 It is unspecified whether a thread that has exited but remains unjoined counts against  
 33436 \_PTHREAD\_THREADS\_MAX.

33437 **RETURN VALUE**

33438 If successful, the *pthread\_join()* function shall return zero; otherwise, an error number shall be  
 33439 returned to indicate the error.

33440 **ERRORS**33441 The *pthread\_join()* function shall fail if:

33442 [EINVAL] The implementation has detected that the value specified by *thread* does not  
 33443 refer to a joinable thread.

33444 [ESRCH] No thread could be found corresponding to that specified by the given thread  
 33445 ID.

33446 The *pthread\_join()* function may fail if:

33447 [EDEADLK] A deadlock was detected or the value of *thread* specifies the calling thread.

33448 The *pthread\_join()* function shall not return an error code of [EINTR].

33449 **EXAMPLES**

33450 An example of thread creation and deletion follows:

```

33451 typedef struct {
33452 int *ar;
33453 long n;
33454 } subarray;

33455 void *
33456 incer(void *arg)
33457 {
33458 long i;

33459 for (i = 0; i < ((subarray *)arg)->n; i++)
33460 ((subarray *)arg)->ar[i]++;
33461 }

33462 main()
33463 {
33464 int ar[1000000];

```

```

33465 pthread_t th1, th2;
33466 subarray sb1, sb2;

33467 sb1.ar = &ar[0];
33468 sb1.n = 500000;
33469 (void) pthread_create(&th1, NULL, incer, &sb1);

33470 sb2.ar = &ar[500000];
33471 sb2.n = 500000;
33472 (void) pthread_create(&th2, NULL, incer, &sb2);

33473 (void) pthread_join(th1, NULL);
33474 (void) pthread_join(th2, NULL);
33475 }
```

**33476 APPLICATION USAGE**

33477 None.

**33478 RATIONALE**

33479 The *pthread\_join()* function is a convenience that has proven useful in multi-threaded  
33480 applications. It is true that a programmer could simulate this function if it were not provided by  
33481 passing extra state as part of the argument to the *start\_routine()*. The terminating thread would  
33482 set a flag to indicate termination and broadcast a condition that is part of that state; a joining  
33483 thread would wait on that condition variable. While such a technique would allow a thread to  
33484 wait on more complex conditions (for example, waiting for multiple threads to terminate),  
33485 waiting on individual thread termination is considered widely useful. Also, including the  
33486 *pthread\_join()* function in no way precludes a programmer from coding such complex waits.  
33487 Thus, while not a primitive, including *pthread\_join()* in this volume of IEEE Std. 1003.1-200x was  
33488 considered valuable.

33489 The *pthread\_join()* function provides a simple mechanism allowing an application to wait for a  
33490 thread to terminate. After the thread terminates, the application may then choose to clean up  
33491 resources that were used by the thread. For instance, after *pthread\_join()* returns, any  
33492 application-provided stack storage could be reclaimed.

33493 The *pthread\_join()* or *pthread\_detach()* function should eventually be called for every thread that  
33494 is created with the *detachstate* attribute set to *PTHREAD\_CREATE\_JOINABLE* so that storage  
33495 associated with the thread may be reclaimed.

33496 The interaction between *pthread\_join()* and cancelation is well-defined for the following reasons:

- 33497 • The *pthread\_join()* function, like all other non-async-cancel-safe functions, can only be called  
33498 with deferred cancelability type.
- 33499 • Cancelation cannot occur in the disabled cancelability state.

33500 Thus, only the default cancelability state need be considered. As specified, either the  
33501 *pthread\_join()* call is canceled, or it succeeds, but not both. The difference is obvious to the  
33502 application, since either a cancelation handler is run or *pthread\_join()* returns. There are no race  
33503 conditions since *pthread\_join()* was called in the deferred cancelability state.

**33504 FUTURE DIRECTIONS**

33505 None.

**33506 SEE ALSO**

33507 *pthread\_create()*, *wait()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

33508 **CHANGE HISTORY**

33509 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33510 **Issue 6**

33511 The *pthread\_join()* function is marked as part of the Threads option.

33512 **NAME**

33513 pthread\_key\_create — thread-specific data key creation

33514 **SYNOPSIS**

33515 THR #include &lt;pthread.h&gt;

33516 int pthread\_key\_create(pthread\_key\_t \*key, void (\*destructor)(void\*));

33517

33518 **DESCRIPTION**

33519 The *pthread\_key\_create()* function shall create a thread-specific data key visible to all threads in  
33520 the process. Key values provided by *pthread\_key\_create()* are opaque objects used to locate  
33521 thread-specific data. Although the same key value may be used by different threads, the values  
33522 bound to the key by *pthread\_setspecific()* are maintained on a per-thread basis and persist for the  
33523 life of the calling thread.

33524 Upon key creation, the value NULL shall be associated with the new key in all active threads.  
33525 Upon thread creation, the value NULL shall be associated with all defined keys in the new  
33526 thread.

33527 An optional destructor function may be associated with each key value. At thread exit, if a key  
33528 value has a non-NULL destructor pointer, and the thread has a non-NULL value associated with  
33529 that key, the value of the key is set to NULL, and then the function pointed to is called with the  
33530 previously associated value as its sole argument. The order of destructor calls is unspecified if  
33531 more than one destructor exists for a thread when it exits.

33532 If, after all the destructors have been called for all non-NULL values with associated destructors,  
33533 there are still some non-NULL values with associated destructors, then the process is repeated.  
33534 If, after at least {PTHREAD\_DESTRUCTOR\_ITERATIONS} iterations of destructor calls for  
33535 outstanding non-NULL values, there are still some non-NULL values with associated  
33536 destructors, implementations may stop calling destructors, or they may continue calling  
33537 destructors until no non-NULL values with associated destructors exist, even though this might  
33538 result in an infinite loop.

33539 **RETURN VALUE**

33540 If successful, the *pthread\_key\_create()* function shall store the newly created key value at *\*key* and  
33541 shall return zero. Otherwise, an error number shall be returned to indicate the error.

33542 **ERRORS**33543 The *pthread\_key\_create()* function shall fail if:

33544 [EAGAIN] The system lacked the necessary resources to create another thread-specific  
33545 data key, or the system-imposed limit on the total number of keys per process  
33546 PTHREAD\_KEYS\_MAX has been exceeded.

33547 [ENOMEM] Insufficient memory exists to create the key.

33548 The *pthread\_key\_create()* function shall not return an error code of [EINTR].

33549 **EXAMPLES**

33550 The following example demonstrates a function that initializes a thread-specific data key when  
 33551 it is first called, and associates a thread-specific object with each calling thread, initializing this  
 33552 object when necessary.

```

33553 static pthread_key_t key;
33554 static pthread_once_t key_once = PTHREAD_ONCE_INIT;

33555 static void
33556 make_key()
33557 {
33558 (void) pthread_key_create(&key, NULL);
33559 }

33560 func()
33561 {
33562 void *ptr;

33563 (void) pthread_once(&key_once, make_key);
33564 if ((ptr = pthread_getspecific(key)) == NULL) {
33565 ptr = malloc(OBJECT_SIZE);
33566 ...
33567 (void) pthread_setspecific(key, ptr);
33568 }
33569 ...
33570 }

```

33571 Note that the key has to be initialized before *pthread\_getspecific()* or *pthread\_setspecific()* can be  
 33572 used. The *pthread\_key\_create()* call could either be explicitly made in a module initialization  
 33573 routine, or it can be done implicitly by the first call to a module as in this example. Any attempt  
 33574 to use the key before it is initialized is a programming error, making the code below incorrect.

```

33575 static pthread_key_t key;

33576 func()
33577 {
33578 void *ptr;

33579 /* KEY NOT INITIALIZED!!! THIS WON'T WORK!!! */
33580 if ((ptr = pthread_getspecific(key)) == NULL &&
33581 pthread_setspecific(key, NULL) != 0) {
33582 pthread_key_create(&key, NULL);
33583 ...
33584 }
33585 }

```

33586 **APPLICATION USAGE**

33587 None.

## 33588 RATIONALE

33589 **Destructor Functions**

33590 Normally, the value bound to a key on behalf of a particular thread is a pointer to storage  
33591 allocated dynamically on behalf of the calling thread. The destructor functions specified with  
33592 *pthread\_key\_create()* are intended to be used to free this storage when the thread exits. Thread  
33593 cancelation cleanup handlers cannot be used for this purpose because thread-specific data may  
33594 persist outside the lexical scope in which the cancelation cleanup handlers operate.

33595 If the value associated with a key needs to be updated during the lifetime of the thread, it may  
33596 be necessary to release the storage associated with the old value before the new value is bound.  
33597 Although the *pthread\_setspecific()* function could do this automatically, this feature is not needed  
33598 often enough to justify the added complexity. Instead, the programmer is responsible for freeing  
33599 the stale storage:

```
33600 pthread_getspecific(key, &old);
33601 new = allocate();
33602 destructor(old);
33603 pthread_setspecific(key, new);
```

33604 **Note:** The above example could leak storage if run with asynchronous cancelation enabled.  
33605 No such problems occur in the default cancelation state if no cancelation points  
33606 occur between the get and set.

33607 There is no notion of a destructor-safe function. If an application does not call *pthread\_exit()*  
33608 from a signal handler, or if it blocks any signal whose handler may call *pthread\_exit()* while  
33609 calling async-unsafe functions, all functions may be safely called from destructors.

33610 **Non-Idempotent Data Key Creation**

33611 There were requests to make *pthread\_key\_create()* idempotent with respect to a given *key* address  
33612 parameter. This would allow applications to call *pthread\_key\_create()* multiple times for a given  
33613 *key* address and be guaranteed that only one key would be created. Doing so would require the  
33614 key value to be previously initialized (possibly at compile time) to a known null value and  
33615 would require that implicit mutual-exclusion be performed based on the address and contents of  
33616 the *key* parameter in order to guarantee that exactly one key would be created.

33617 Unfortunately, the implicit mutual-exclusion would not be limited to only *pthread\_key\_create()*.  
33618 On many implementations, implicit mutual-exclusion would also have to be performed by  
33619 *pthread\_getspecific()* and *pthread\_setspecific()* in order to guard against using incompletely stored  
33620 or not-yet-visible key values. This could significantly increase the cost of important operations,  
33621 particularly *pthread\_getspecific()*.

33622 Thus, this proposal was rejected. The *pthread\_key\_create()* function performs no implicit  
33623 synchronization. It is the responsibility of the programmer to ensure that it is called exactly once  
33624 per key before use of the key. Several straightforward mechanisms can already be used to  
33625 accomplish this, including calling explicit module initialization functions, using mutexes, and  
33626 using *pthread\_once()*. This places no significant burden on the programmer, introduces no  
33627 possibly confusing *ad hoc* implicit synchronization mechanism, and potentially allows  
33628 commonly used thread-specific data operations to be more efficient.

33629 **FUTURE DIRECTIONS**

33630 None.

33631 **SEE ALSO**

33632 *pthread\_getspecific()*, *pthread\_key\_delete()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
33633 <pthread.h>

33634 **CHANGE HISTORY**

33635 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33636 **Issue 6**

33637 The *pthread\_key\_create()* function is marked as part of the Threads option. |

33638 IEEE PASC Interpretation 1003.1c #8 is applied, updating the DESCRIPTION. |

33639 **NAME**

33640 pthread\_key\_delete — thread-specific data key deletion

33641 **SYNOPSIS**33642 THR `#include <pthread.h>`33643 `int pthread_key_delete(pthread_key_t key);`

33644

33645 **DESCRIPTION**

33646 The *pthread\_key\_delete()* function shall delete a thread-specific data key previously returned by  
33647 *pthread\_key\_create()*. The thread-specific data values associated with *key* need not be NULL at  
33648 the time *pthread\_key\_delete()* is called. It is the responsibility of the application to free any  
33649 application storage or perform any cleanup actions for data structures related to the deleted key  
33650 or associated thread-specific data in any threads; this cleanup can be done either before or after  
33651 *pthread\_key\_delete()* is called. Any attempt to use *key* following the call to *pthread\_key\_delete()*  
33652 results in undefined behavior.

33653 The *pthread\_key\_delete()* function shall be callable from within destructor functions. No  
33654 destructor functions shall be invoked by *pthread\_key\_delete()*. Any destructor function that may  
33655 have been associated with *key* shall no longer be called upon thread exit.

33656 **RETURN VALUE**

33657 If successful, the *pthread\_key\_delete()* function shall return zero; otherwise, an error number shall  
33658 be returned to indicate the error.

33659 **ERRORS**33660 The *pthread\_key\_delete()* function may fail if:33661 [EINVAL] The *key* value is invalid.33662 The *pthread\_key\_delete()* function shall not return an error code of [EINTR].33663 **EXAMPLES**

33664 None.

33665 **APPLICATION USAGE**

33666 None.

33667 **RATIONALE**

33668 A thread-specific data key deletion function has been included in order to allow the resources  
33669 associated with an unused thread-specific data key to be freed. Unused thread-specific data keys  
33670 can arise, among other scenarios, when a dynamically loaded module that allocated a key is  
33671 unloaded.

33672 Portable applications are responsible for performing any cleanup actions needed for data  
33673 structures associated with the key to be deleted, including data referenced by thread-specific  
33674 data values. No such cleanup is done by *pthread\_key\_delete()*. In particular, destructor functions  
33675 are not called. There are several reasons for this division of responsibility:

- 33676 1. The associated destructor functions used to free thread-specific data at thread exit time are  
33677 only guaranteed to work correctly when called in the thread that allocated the thread-  
33678 specific data. (Destructors themselves may utilize thread-specific data.) Thus, they cannot  
33679 be used to free thread-specific data in other threads at key deletion time. Attempting to  
33680 have them called by other threads at key deletion time would require other threads to be  
33681 asynchronously interrupted. But since interrupted threads could be in an arbitrary state,  
33682 including holding locks necessary for the destructor to run, this approach would fail. In  
33683 general, there is no safe mechanism whereby an implementation could free thread-specific  
33684 data at key deletion time.

33685           2. Even if there were a means of safely freeing thread-specific data associated with keys to be  
33686 deleted, doing so would require that implementations be able to enumerate the threads  
33687 with non-NULL data and potentially keep them from creating more thread-specific data  
33688 while the key deletion is occurring. This special case could cause extra synchronization in  
33689 the normal case, which would otherwise be unnecessary.

33690           For an application to know that it is safe to delete a key, it has to know that all the threads that  
33691 might potentially ever use the key do not attempt to use it again. For example, it could know this  
33692 if all the client threads have called a cleanup procedure declaring that they are through with the  
33693 module that is being shut down, perhaps by zero'ing a reference count.

33694 **FUTURE DIRECTIONS**

33695           None.

33696 **SEE ALSO**

33697           *pthread\_key\_create()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

33698 **CHANGE HISTORY**

33699           First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33700 **Issue 6**

33701           The *pthread\_key\_delete()* function is marked as part of the Threads option.

33702 **NAME**

33703 pthread\_kill — send a signal to a thread

33704 **SYNOPSIS**

33705 THR #include &lt;signal.h&gt;

33706 int pthread\_kill(pthread\_t thread, int sig);

33707

33708 **DESCRIPTION**33709 The *pthread\_kill()* function is used to request that a signal be delivered to the specified thread.33710 As in *kill()*, if *sig* is zero, error checking is performed but no signal is actually sent.33711 **RETURN VALUE**

33712 Upon successful completion, the function shall return a value of zero. Otherwise, the function

33713 shall return an error number. If the *pthread\_kill()* function fails, no signal shall be sent.33714 **ERRORS**33715 The *pthread\_kill()* function shall fail if:

33716 [ESRCH] No thread could be found corresponding to that specified by the given thread ID. |

33717

33718 [EINVAL] The value of the *sig* argument is an invalid or unsupported signal number. |33719 The *pthread\_kill()* function shall not return an error code of [EINTR]. |33720 **EXAMPLES**

33721 None.

33722 **APPLICATION USAGE**33723 The *pthread\_kill()* function provides a mechanism for asynchronously directing a signal at a  
33724 thread in the calling process. This could be used, for example, by one thread to affect broadcast  
33725 delivery of a signal to a set of threads.33726 Note that *pthread\_kill()* only causes the signal to be handled in the context of the given thread;  
33727 the signal action (termination or stopping) affects the process as a whole.33728 **RATIONALE**

33729 None.

33730 **FUTURE DIRECTIONS**

33731 None.

33732 **SEE ALSO**33733 *kill()*, *pthread\_self()*, *raise()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <signal.h> |33734 **CHANGE HISTORY**

33735 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33736 **Issue 6**33737 The *pthread\_kill()* function is marked as part of the Threads option. |

33738 The APPLICATION USAGE section is added.

33739 **NAME**

33740 pthread\_mutex\_destroy, pthread\_mutex\_init — destroy and initialize a mutex

33741 **SYNOPSIS**

33742 THR #include &lt;pthread.h&gt;

```

33743 int pthread_mutex_destroy(pthread_mutex_t *mutex);
33744 int pthread_mutex_init(pthread_mutex_t *restrict mutex,
33745 const pthread_mutexattr_t *restrict attr);
33746 pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
33747

```

33748 **DESCRIPTION**

33749 The *pthread\_mutex\_destroy()* function destroys the mutex object referenced by *mutex*; the mutex  
 33750 object becomes, in effect, uninitialized. An implementation may cause *pthread\_mutex\_destroy()* to  
 33751 set the object referenced by *mutex* to an invalid value. A destroyed mutex object can be re-  
 33752 initialized using *pthread\_mutex\_init()*; the results of otherwise referencing the object after it has  
 33753 been destroyed are undefined.

33754 It shall be safe to destroy an initialized mutex that is unlocked. Attempting to destroy a locked  
 33755 mutex results in undefined behavior.

33756 The *pthread\_mutex\_init()* function initializes the mutex referenced by *mutex* with attributes  
 33757 specified by *attr*. If *attr* is NULL, the default mutex attributes are used; the effect shall be the  
 33758 same as passing the address of a default mutex attributes object. Upon successful initialization,  
 33759 the state of the mutex becomes initialized and unlocked.

33760 Only *mutex* itself may be used for performing synchronization. The result of referring to copies  
 33761 of *mutex* in calls to *pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, *pthread\_mutex\_unlock()*, and  
 33762 *pthread\_mutex\_destroy()* is undefined.

33763 Attempting to initialize an already initialized mutex results in undefined behavior.

33764 In cases where default mutex attributes are appropriate, the macro  
 33765 PTHREAD\_MUTEX\_INITIALIZER can be used to initialize mutexes that are statically allocated.  
 33766 The effect shall be equivalent to dynamic initialization by a call to *pthread\_mutex\_init()* with  
 33767 parameter *attr* specified as NULL, except that no error checks are performed.

33768 **RETURN VALUE**

33769 If successful, the *pthread\_mutex\_destroy()* and *pthread\_mutex\_init()* functions shall return zero;  
 33770 otherwise, an error number shall be returned to indicate the error.

33771 The [EBUSY] and [EINVAL] error checks, if implemented, act as if they were performed  
 33772 immediately at the beginning of processing for the function and shall cause an error return prior  
 33773 to modifying the state of the mutex specified by *mutex*.

33774 **ERRORS**

33775 The *pthread\_mutex\_destroy()* function may fail if:

33776 [EBUSY] The implementation has detected an attempt to destroy the object referenced  
 33777 by *mutex* while it is locked or referenced (for example, while being used in a  
 33778 *pthread\_cond\_timedwait()* or *pthread\_cond\_wait()*) by another thread.

33779 [EINVAL] The value specified by *mutex* is invalid.

33780 The *pthread\_mutex\_init()* function shall fail if:

33781 [EAGAIN] The system lacked the necessary resources (other than memory) to initialize  
 33782 another mutex.

- 33783 [ENOMEM] Insufficient memory exists to initialize the mutex.
- 33784 [EPERM] The caller does not have the privilege to perform the operation.
- 33785 The *pthread\_mutex\_init()* function may fail if:
- 33786 [EBUSY] The implementation has detected an attempt to re-initialize the object  
33787 referenced by *mutex*, a previously initialized, but not yet destroyed, mutex.
- 33788 [EINVAL] The value specified by *attr* is invalid.
- 33789 These functions shall not return an error code of [EINTR].

**33790 EXAMPLES**

33791 None.

**33792 APPLICATION USAGE**

33793 None.

**33794 RATIONALE****33795 Alternate Implementations Possible**

33796 This volume of IEEE Std. 1003.1-200x supports several alternative implementations of mutexes.  
33797 An implementation may store the lock directly in the object of type **pthread\_mutex\_t**.  
33798 Alternatively, an implementation may store the lock in the heap and merely store a pointer,  
33799 handle, or unique ID in the mutex object. Either implementation has advantages or may be  
33800 required on certain hardware configurations. So that portable code can be written that is  
33801 invariant to this choice, this volume of IEEE Std. 1003.1-200x does not define assignment or  
33802 equality for this type, and it uses the term “initialize” to reinforce the (more restrictive) notion  
33803 that the lock may actually reside in the mutex object itself.

33804 Note that this precludes an over-specification of the type of the mutex or condition variable and  
33805 motivates the opacity of the type.

33806 An implementation is permitted, but not required, to have *pthread\_mutex\_destroy()* store an  
33807 illegal value into the mutex. This may help detect erroneous programs that try to lock (or  
33808 otherwise reference) a mutex that has already been destroyed.

**33809 Tradeoff Between Error Checks and Performance Supported**

33810 Many of the error checks were made optional in order to let implementations trade off  
33811 performance *versus* degree of error checking according to the needs of their specific applications  
33812 and execution environment. As a general rule, errors or conditions caused by the system (such as  
33813 insufficient memory) always need to be reported, but errors due to an erroneously coded  
33814 application (such as failing to provide adequate synchronization to prevent a mutex from being  
33815 deleted while in use) are made optional.

33816 A wide range of implementations is thus made possible. For example, an implementation  
33817 intended for application debugging may implement all of the error checks, but an  
33818 implementation running a single, provably correct application under very tight performance  
33819 constraints in an embedded computer might implement minimal checks. An implementation  
33820 might even be provided in two versions, similar to the options that compilers provide: a full-  
33821 checking, but slower version; and a limited-checking, but faster version. To forbid this  
33822 optionality would be a disservice to users.

33823 By carefully limiting the use of “undefined behavior” only to things that an erroneous (badly  
33824 coded) application might do, and by defining that resource-not-available errors are mandatory,  
33825 this volume of IEEE Std. 1003.1-200x ensures that a fully-conforming application is portable

33826 across the full range of implementations, while not forcing all implementations to add overhead  
33827 to check for numerous things that a correct program never does.

### 33828 **Why No Limits Defined**

33829 Defining symbols for the maximum number of mutexes and condition variables was considered  
33830 but rejected because the number of these objects may change dynamically. Furthermore, many  
33831 implementations place these objects into application memory; thus, there is no explicit  
33832 maximum.

### 33833 **Static Initializers for Mutexes and Condition Variables**

33834 Providing for static initialization of statically allocated synchronization objects allows modules  
33835 with private static synchronization variables to avoid runtime initialization tests and overhead.  
33836 Furthermore, it simplifies the coding of self-initializing modules. Such modules are common in  
33837 C libraries, where for various reasons the design calls for self-initialization instead of requiring  
33838 an explicit module initialization function to be called. An example use of static initialization  
33839 follows.

33840 Without static initialization, a self-initializing routine *foo()* might look as follows:

```
33841 static pthread_once_t foo_once = PTHREAD_ONCE_INIT;
33842 static pthread_mutex_t foo_mutex;

33843 void foo_init()
33844 {
33845 pthread_mutex_init(&foo_mutex, NULL);
33846 }

33847 void foo()
33848 {
33849 pthread_once(&foo_once, foo_init);
33850 pthread_mutex_lock(&foo_mutex);
33851 /* Do work. */
33852 pthread_mutex_unlock(&foo_mutex);
33853 }
```

33854 With static initialization, the same routine could be coded as follows:

```
33855 static pthread_mutex_t foo_mutex = PTHREAD_MUTEX_INITIALIZER;

33856 void foo()
33857 {
33858 pthread_mutex_lock(&foo_mutex);
33859 /* Do work. */
33860 pthread_mutex_unlock(&foo_mutex);
33861 }
```

33862 Note that the static initialization both eliminates the need for the initialization test inside  
33863 *pthread\_once()* and the fetch of *&foo\_mutex* to learn the address to be passed to  
33864 *pthread\_mutex\_lock()* or *pthread\_mutex\_unlock()*.

33865 Thus, the C code written to initialize static objects is simpler on all systems and is also faster on a  
33866 large class of systems; those where the (entire) synchronization object can be stored in  
33867 application memory.

33868 Yet the locking performance question is likely to be raised for machines that require mutexes to  
33869 be allocated out of special memory. Such machines actually have to have mutexes and possibly

33870 condition variables contain pointers to the actual hardware locks. For static initialization to work  
33871 on such machines, *pthread\_mutex\_lock()* also has to test whether or not the pointer to the actual  
33872 lock has been allocated. If it has not, *pthread\_mutex\_lock()* has to initialize it before use. The  
33873 reservation of such resources can be made when the program is loaded, and hence return codes  
33874 have not been added to mutex locking and condition variable waiting to indicate failure to  
33875 complete initialization.

33876 This runtime test in *pthread\_mutex\_lock()* would at first seem to be extra work; an extra test is  
33877 required to see whether the pointer has been initialized. On most machines this would actually  
33878 be implemented as a fetch of the pointer, testing the pointer against zero, and then using the  
33879 pointer if it has already been initialized. While the test might seem to add extra work, the extra  
33880 effort of testing a register is usually negligible since no extra memory references are actually  
33881 done. As more and more machines provide caches, the real expenses are memory references, not  
33882 instructions executed.

33883 Alternatively, depending on the machine architecture, there are often ways to eliminate *all*  
33884 overhead in the most important case: on the lock operations that occur *after* the lock has been  
33885 initialized. This can be done by shifting more overhead to the less frequent operation:  
33886 initialization. Since out-of-line mutex allocation also means that an address has to be  
33887 dereferenced to find the actual lock, one technique that is widely applicable is to have static  
33888 initialization store a bogus value for that address; in particular, an address that causes a machine  
33889 fault to occur. When such a fault occurs upon the first attempt to lock such a mutex, validity  
33890 checks can be done, and then the correct address for the actual lock can be filled in. Subsequent  
33891 lock operations incur no extra overhead since they do not “fault”. This is merely one technique  
33892 that can be used to support static initialization, while not adversely affecting the performance of  
33893 lock acquisition. No doubt there are other techniques that are highly machine-dependent.

33894 The locking overhead for machines doing out-of-line mutex allocation is thus similar for  
33895 modules being implicitly initialized, where it is improved for those doing mutex allocation  
33896 entirely inline. The inline case is thus made much faster, and the out-of-line case is not  
33897 significantly worse.

33898 Besides the issue of locking performance for such machines, a concern is raised that it is possible  
33899 that threads would serialize contending for initialization locks when attempting to finish  
33900 initializing statically allocated mutexes. (Such finishing would typically involve taking an  
33901 internal lock, allocating a structure, storing a pointer to the structure in the mutex, and releasing  
33902 the internal lock.) First, many implementations would reduce such serialization by hashing on  
33903 the mutex address. Second, such serialization can only occur a bounded number of times. In  
33904 particular, it can happen at most as many times as there are statically allocated synchronization  
33905 objects. Dynamically allocated objects would still be initialized via *pthread\_mutex\_init()* or  
33906 *pthread\_cond\_init()*.

33907 Finally, if none of the above optimization techniques for out-of-line allocation yields sufficient  
33908 performance for an application on some implementation, the application can avoid static  
33909 initialization altogether by explicitly initializing all synchronization objects with the  
33910 corresponding *pthread\_\*\_init()* functions, which are supported by all implementations. An  
33911 implementation can also document the tradeoffs and advise which initialization technique is  
33912 more efficient for that particular implementation.

33913 **Destroying Mutexes**

33914 A mutex can be destroyed immediately after it is unlocked. For example, consider the following  
33915 code:

```
33916 struct obj {
33917 pthread_mutex_t om;
33918 int refcnt;
33919 ...
33920 };

33921 obj_done(struct obj *op)
33922 {
33923 pthread_mutex_lock(&op->om);
33924 if (--op->refcnt == 0) {
33925 pthread_mutex_unlock(&op->om);
33926 (A) pthread_mutex_destroy(&op->om);
33927 (B) free(op);
33928 } else
33929 (C) pthread_mutex_unlock(&op->om);
33930 }
```

33931 In this case *obj* is reference counted and *obj\_done()* is called whenever a reference to the object is  
33932 dropped. Implementations are required to allow an object to be destroyed and freed and  
33933 potentially unmapped (for example, lines A and B) immediately after the object is unlocked (line  
33934 C).

33935 **FUTURE DIRECTIONS**

33936 None.

33937 **SEE ALSO**

33938 *pthread\_mutex\_getprioceiling()*, *pthread\_mutex\_lock()*, *pthread\_mutex\_timedlock()*,  
33939 *pthread\_mutexattr\_getpshared()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
33940 <pthread.h>

33941 **CHANGE HISTORY**

33942 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33943 **Issue 6**

33944 The *pthread\_mutex\_destroy()* and *pthread\_mutex\_init()* functions are marked as part of the  
33945 Threads option.

33946 The *pthread\_mutex\_timedlock()* function is added to the SEE ALSO section for alignment with  
33947 IEEE Std. 1003.1d-1999.

33948 IEEE PASC Interpretation 1003.1c #34 is applied, updating the DESCRIPTION.

33949 The **restrict** keyword is added to the *pthread\_mutex\_init()* prototype for alignment with the  
33950 ISO/IEC 9899:1999 standard.

33951 **NAME**

33952 pthread\_mutex\_getprioceiling, pthread\_mutex\_setprioceiling — change the priority ceiling of a  
 33953 mutex (**REALTIME THREADS**)

33954 **SYNOPSIS**

```
33955 TPP #include <pthread.h>
33956 int pthread_mutex_getprioceiling(const pthread_mutex_t *restrict mutex,
33957 int *restrict prioceiling);
33958 int pthread_mutex_setprioceiling(pthread_mutex_t *restrict mutex,
33959 int prioceiling, int *restrict old_ceiling);
33960
```

33961 **DESCRIPTION**

33962 The *pthread\_mutex\_getprioceiling()* function shall return the current priority ceiling of the mutex.

33963 The *pthread\_mutex\_setprioceiling()* function either locks the mutex if it is unlocked, or blocks until  
 33964 it can successfully lock the mutex, then it changes the mutex's priority ceiling and releases the  
 33965 mutex. When the change is successful, the previous value of the priority ceiling is returned in  
 33966 *old\_ceiling*. The process of locking the mutex need not adhere to the priority protect protocol.

33967 If the *pthread\_mutex\_setprioceiling()* function fails, the mutex priority ceiling shall not be  
 33968 changed.

33969 **RETURN VALUE**

33970 If successful, the *pthread\_mutex\_getprioceiling()* and *pthread\_mutex\_setprioceiling()* functions shall  
 33971 return zero; otherwise, an error number shall be returned to indicate the error.

33972 **ERRORS**

33973 The *pthread\_mutex\_getprioceiling()* and *pthread\_mutex\_setprioceiling()* functions may fail if:

33974 [EINVAL] The priority requested by *prioceiling* is out of range.

33975 [EINVAL] The value specified by *mutex* does not refer to a currently existing mutex.

33976 [EPERM] The caller does not have the privilege to perform the operation.

33977 These functions shall not return an error code of [EINTR].

33978 **EXAMPLES**

33979 None.

33980 **APPLICATION USAGE**

33981 None.

33982 **RATIONALE**

33983 None.

33984 **FUTURE DIRECTIONS**

33985 None.

33986 **SEE ALSO**

33987 *pthread\_mutex\_destroy()*, *pthread\_mutex\_lock()*, *pthread\_mutex\_timedlock()*, the Base Definitions  
 33988 volume of IEEE Std. 1003.1-200x, <pthread.h>

33989 **CHANGE HISTORY**

33990 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

33991 Marked as part of the Realtime Threads Feature Group.

33992 **Issue 6**

33993 The *pthread\_mutex\_getprioceiling()* and *pthread\_mutex\_setprioceiling()* functions are marked as  
33994 part of the Thread Priority Protection option.

33995 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
33996 implementation does not support the Thread Priority Protection option.

33997 The [ENOSYS] error denoting non-support of the priority ceiling protocol for mutexes has been  
33998 removed. This is since if the implementation provides the functions (regardless of whether  
33999 `_POSIX_PTHREAD_PRIO_PROTECT` is defined), they must function as in the DESCRIPTION  
34000 and therefore the priority ceiling protocol for mutexes is supported.

34001 The *pthread\_mutex\_timedlock()* function is added to the SEE ALSO section for alignment with  
34002 IEEE Std. 1003.1d-1999.

34003 The **restrict** keyword is added to the *pthread\_mutex\_getprioceiling()* and  
34004 *pthread\_mutex\_setprioceiling()* prototypes for alignment with the ISO/IEC 9899:1999 standard.

34005 **NAME**

34006 pthread\_mutex\_init — initialize a mutex

34007 **SYNOPSIS**

34008 THR #include <pthread.h>

34009 int pthread\_mutex\_init(pthread\_mutex\_t \*restrict mutex,

34010 const pthread\_mutexattr\_t \*restrict attr);

34011 pthread\_mutex\_t mutex = PTHREAD\_MUTEX\_INITIALIZER;

34012

34013 **DESCRIPTION**

34014 Refer to *pthread\_mutex\_destroy()*.

34015 **NAME**

34016 pthread\_mutex\_lock, pthread\_mutex\_trylock, pthread\_mutex\_unlock — lock and unlock a  
34017 mutex

34018 **SYNOPSIS**

34019 THR #include <pthread.h>

34020 int pthread\_mutex\_lock(pthread\_mutex\_t \*mutex);

34021 int pthread\_mutex\_trylock(pthread\_mutex\_t \*mutex);

34022 int pthread\_mutex\_unlock(pthread\_mutex\_t \*mutex);

34023

34024 **DESCRIPTION**

34025 The mutex object referenced by *mutex* shall be locked by calling *pthread\_mutex\_lock()*. If the  
34026 mutex is already locked, the calling thread shall block until the mutex becomes available. This  
34027 operation shall return with the mutex object referenced by *mutex* in the locked state with the  
34028 calling thread as its owner.

34029 XSI If the mutex type is PTHREAD\_MUTEX\_NORMAL, deadlock detection shall not be provided.  
34030 Attempting to relock the mutex causes deadlock. If a thread attempts to unlock a mutex that it  
34031 has not locked or a mutex which is unlocked, undefined behavior results.

34032 If the mutex type is PTHREAD\_MUTEX\_ERRORCHECK, then error checking shall be provided.  
34033 If a thread attempts to relock a mutex that it has already locked, an error is returned. If a thread  
34034 attempts to unlock a mutex that it has not locked or a mutex which is unlocked, an error is  
34035 returned.

34036 If the mutex type is PTHREAD\_MUTEX\_RECURSIVE, then the mutex shall maintain the  
34037 concept of a lock count. When a thread successfully acquires a mutex for the first time, the lock  
34038 count is set to one. Every time a thread relocks this mutex, the lock count is incremented by one.  
34039 Each time the thread unlocks the mutex, the lock count is decremented by one. When the lock  
34040 count reaches zero, the mutex becomes available for other threads to acquire. If a thread  
34041 attempts to unlock a mutex that it has not locked or a mutex which is unlocked, an error is  
34042 returned.

34043 If the mutex type is PTHREAD\_MUTEX\_DEFAULT, attempting to recursively lock the mutex  
34044 results in undefined behavior. Attempting to unlock the mutex if it was not locked by the calling  
34045 thread results in undefined behavior. Attempting to unlock the mutex if it is not locked results in  
34046 undefined behavior.

34047 The *pthread\_mutex\_trylock()* function is identical to *pthread\_mutex\_lock()* except that if the mutex  
34048 object referenced by *mutex* is currently locked (by any thread, including the current thread), the  
34049 call shall return immediately.

34050 XSI The *pthread\_mutex\_unlock()* function releases the mutex object referenced by *mutex*. The manner  
34051 in which a mutex is released is dependent upon the mutex's type attribute. If there are threads  
34052 blocked on the mutex object referenced by *mutex* when *pthread\_mutex\_unlock()* is called,  
34053 resulting in the mutex becoming available, the scheduling policy is used to determine which  
34054 thread shall acquire the mutex.

34055 XSI (In the case of PTHREAD\_MUTEX\_RECURSIVE mutexes, the mutex shall become available  
34056 when the count reaches zero and the calling thread no longer has any locks on this mutex).

34057 If a signal is delivered to a thread waiting for a mutex, upon return from the signal handler the  
34058 thread shall resume waiting for the mutex as if it was not interrupted.

34059 **RETURN VALUE**

34060 If successful, the *pthread\_mutex\_lock()* and *pthread\_mutex\_unlock()* functions shall return zero;  
34061 otherwise, an error number shall be returned to indicate the error.

34062 The *pthread\_mutex\_trylock()* function shall return zero if a lock on the mutex object referenced by  
34063 *mutex* is acquired. Otherwise, an error number is returned to indicate the error.

34064 **ERRORS**

34065 The *pthread\_mutex\_lock()* and *pthread\_mutex\_trylock()* functions shall fail if:

34066 [EINVAL] The *mutex* was created with the protocol attribute having the value  
34067 PTHREAD\_PRIO\_PROTECT and the calling thread's priority is higher than  
34068 the mutex's current priority ceiling.

34069 The *pthread\_mutex\_trylock()* function shall fail if:

34070 [EBUSY] The *mutex* could not be acquired because it was already locked.

34071 The *pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, and *pthread\_mutex\_unlock()* functions may  
34072 fail if:

34073 [EINVAL] The value specified by *mutex* does not refer to an initialized mutex object.

34074 XSI [EAGAIN] The mutex could not be acquired because the maximum number of recursive  
34075 locks for *mutex* has been exceeded.

34076 The *pthread\_mutex\_lock()* function may fail if:

34077 [EDEADLK] The current thread already owns the mutex.

34078 The *pthread\_mutex\_unlock()* function may fail if:

34079 [EPERM] The current thread does not own the mutex.

34080 These functions shall not return an error code of [EINTR].

34081 **EXAMPLES**

34082 None.

34083 **APPLICATION USAGE**

34084 None.

34085 **RATIONALE**

34086 Mutex objects are intended to serve as a low-level primitive from which other thread  
34087 synchronization functions can be built. As such, the implementation of mutexes should be as  
34088 efficient as possible, and this has ramifications on the features available at the interface.

34089 The mutex functions and the particular default settings of the mutex attributes have been  
34090 motivated by the desire to not preclude fast, inlined implementations of mutex locking and  
34091 unlocking.

34092 For example, deadlocking on a double-lock is explicitly allowed behavior in order to avoid  
34093 requiring more overhead in the basic mechanism than is absolutely necessary. (More "friendly"  
34094 mutexes that detect deadlock or that allow multiple locking by the same thread are easily  
34095 constructed by the user via the other mechanisms provided. For example, *pthread\_self()* can be  
34096 used to record mutex ownership.) Implementations might also choose to provide such extended  
34097 features as options via special mutex attributes.

34098 Since most attributes only need to be checked when a thread is going to be blocked, the use of  
34099 attributes does not slow the (common) mutex-locking case.

34100 Likewise, while being able to extract the thread ID of the owner of a mutex might be desirable, it  
34101 would require storing the current thread ID when each mutex is locked, and this could incur  
34102 unacceptable levels of overhead. Similar arguments apply to a *mutex\_tryunlock* operation.

34103 **FUTURE DIRECTIONS**

34104 None.

34105 **SEE ALSO**

34106 *pthread\_mutex\_destroy()*, *pthread\_mutex\_timedlock()*, the Base Definitions volume of  
34107 IEEE Std. 1003.1-200x, <pthread.h>

34108 **CHANGE HISTORY**

34109 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34110 **Issue 6**

34111 The *pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, and *pthread\_mutex\_unlock()* functions are  
34112 marked as part of the Threads option.

34113 The following new requirements on POSIX implementations derive from alignment with the  
34114 Single UNIX Specification:

- 34115 • The behavior when attempting to relock a mutex is defined.

34116 The *pthread\_mutex\_timedlock()* function is added to the SEE ALSO section for alignment with  
34117 IEEE Std. 1003.1d-1999.

34118 **NAME**

34119 pthread\_mutex\_setprioceiling — change the priority ceiling of a mutex (**REALTIME**  
34120 **THREADS**)

34121 **SYNOPSIS**

34122 TPP #include <pthread.h>

34123 int pthread\_mutex\_setprioceiling(pthread\_mutex\_t \*restrict mutex,  
34124 int prioceiling, int \*restrict old\_ceiling);

34125

34126 **DESCRIPTION**

34127 Refer to *pthread\_mutex\_getprioceiling()*.

34128 **NAME**34129 pthread\_mutex\_timedlock — lock a mutex (**REALTIME THREADS**)34130 **SYNOPSIS**

34131 THR TMO #include &lt;pthread.h&gt;

34132 #include &lt;time.h&gt;

34133 int pthread\_mutex\_timedlock(pthread\_mutex\_t \*restrict mutex,

34134 const struct timespec \*restrict abs\_timeout);

34135

34136 **DESCRIPTION**

34137 The *pthread\_mutex\_timedlock()* function is called to lock the mutex object referenced by *mutex*. If  
 34138 the mutex is already locked, the calling thread blocks until the mutex becomes available as in the  
 34139 *pthread\_mutex\_lock()* function. If the mutex cannot be locked without waiting for another thread  
 34140 to unlock the mutex, this wait shall be terminated when the specified timeout expires.

34141 The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the  
 34142 clock on which timeouts are based (that is, when the value of that clock equals or exceeds  
 34143 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time  
 34144 of the call. If the Timers option is supported, the timeout is based on the `CLOCK_REALTIME`  
 34145 clock; if the Timers option is not supported, the timeout is based on the system clock as returned  
 34146 by the *time()* function. The resolution of the timeout is the resolution of the clock on which it is  
 34147 based. The `timespec` datatype is defined as a structure in the `<time.h>` header.

34148 Under no circumstance will the function fail with a timeout if the mutex can be locked  
 34149 immediately. The validity of the *abs\_timeout* parameter need not be checked if the mutex can be  
 34150 locked immediately.

34151 As a consequence of the priority inheritance rules (for mutexes initialized with the  
 34152 `PRIO_INHERIT` protocol), if a timed mutex wait is terminated because its timeout expires, the  
 34153 priority of the owner of the mutex will be adjusted as necessary to reflect the fact that this thread  
 34154 is no longer among the threads waiting for the mutex.

34155 **RETURN VALUE**

34156 If successful, the *pthread\_mutex\_timedlock()* function shall return zero; otherwise, an error  
 34157 number shall be returned to indicate the error.

34158 **ERRORS**

34159 The *pthread\_mutex\_timedlock()* function shall fail if:

34160 [EINVAL] The mutex was created with the protocol attribute having the value  
 34161 `PTHREAD_PRIO_PROTECT` and the calling thread's priority is higher than  
 34162 the mutex' current priority ceiling.

34163 [EINVAL] The process or thread would have blocked, and the *abs\_timeout* parameter  
 34164 specified a nanoseconds field value less than zero or greater than or equal to  
 34165 1 000 million.

34166 [ETIMEDOUT] The mutex could not be locked before the specified timeout expired.

34167 The *pthread\_mutex\_timedlock()* function may fail if:

34168 [EINVAL] The value specified by *mutex* does not refer to an initialized mutex object.

34169 XSI [EAGAIN] The mutex could not be acquired because the maximum number of recursive  
 34170 locks for mutex has been exceeded.

34171 [EDEADLK] The current thread already owns the mutex.

34172 This function shall not return an error code of [EINTR].

34173 **EXAMPLES**

34174 None.

34175 **APPLICATION USAGE**

34176 The *pthread\_mutex\_timedlock()* function is part of the Threads and Timeouts options and need  
34177 not be provided on all implementations.

34178 **RATIONALE**

34179 None.

34180 **FUTURE DIRECTIONS**

34181 None.

34182 **SEE ALSO**

34183 *pthread\_mutex\_destroy()*, *pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, *time()*, the Base  
34184 Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>, <time.h>

34185 **CHANGE HISTORY**

34186 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

34187 **NAME**

34188 pthread\_mutex\_trylock, pthread\_mutex\_unlock — lock and unlock a mutex

34189 **SYNOPSIS**

34190 THR #include &lt;pthread.h&gt;

34191 int pthread\_mutex\_trylock(pthread\_mutex\_t \*mutex);

34192 int pthread\_mutex\_unlock(pthread\_mutex\_t \*mutex);

34193

34194 **DESCRIPTION**34195 Refer to *pthread\_mutex\_lock()*.

## 34196 NAME

34197 pthread\_mutexattr\_destroy, pthread\_mutexattr\_init — destroy and initialize mutex attributes  
34198 object

## 34199 SYNOPSIS

```
34200 THR #include <pthread.h>
```

```
34201 int pthread_mutexattr_destroy(pthread_mutexattr_t *attr);
```

```
34202 int pthread_mutexattr_init(pthread_mutexattr_t *attr);
```

34203

## 34204 DESCRIPTION

34205 The *pthread\_mutexattr\_destroy()* function destroys a mutex attributes object; the object becomes,  
34206 in effect, uninitialized. An implementation may cause *pthread\_mutexattr\_destroy()* to set the  
34207 object referenced by *attr* to an invalid value. A destroyed mutex attributes object can be re-  
34208 initialized using *pthread\_mutexattr\_init()*; the results of otherwise referencing the object after it  
34209 has been destroyed are undefined.

34210 The *pthread\_mutexattr\_init()* function initializes a mutex attributes object *attr* with the default  
34211 value for all of the attributes defined by the implementation.

34212 The effect of initializing an already initialized mutex attributes object is undefined.

34213 After a mutex attributes object has been used to initialize one or more mutexes, any function  
34214 affecting the attributes object (including destruction) does not affect any previously initialized  
34215 mutexes.

## 34216 RETURN VALUE

34217 Upon successful completion, *pthread\_mutexattr\_destroy()* and *pthread\_mutexattr\_init()* shall  
34218 return zero; otherwise, an error number shall be returned to indicate the error.

## 34219 ERRORS

34220 The *pthread\_mutexattr\_destroy()* function may fail if:

34221 [EINVAL] The value specified by *attr* is invalid.

34222 The *pthread\_mutexattr\_init()* function shall fail if:

34223 [ENOMEM] Insufficient memory exists to initialize the mutex attributes object.

34224 These functions shall not return an error code of [EINTR].

## 34225 EXAMPLES

34226 None.

## 34227 APPLICATION USAGE

34228 None.

## 34229 RATIONALE

34230 See *pthread\_attr\_init()* for a general explanation of attributes. Attributes objects allow  
34231 implementations to experiment with useful extensions and permit extension of this volume of  
34232 IEEE Std. 1003.1-200x without changing the existing functions. Thus, they provide for future  
34233 extensibility of this volume of IEEE Std. 1003.1-200x and reduce the temptation to standardize  
34234 prematurely on semantics that are not yet widely implemented or understood.

34235 Examples of possible additional mutex attributes that have been discussed are *spin\_only*,  
34236 *limited\_spin*, *no\_spin*, *recursive*, and *metered*. (To explain what the latter attributes might mean:  
34237 recursive mutexes would allow for multiple re-locking by the current owner; metered mutexes  
34238 would transparently keep records of queue length, wait time, and so on.) Since there is not yet  
34239 wide agreement on the usefulness of these resulting from shared implementation and usage

34240 experience, they are not yet specified in this volume of IEEE Std. 1003.1-200x. Mutex attributes  
 34241 objects, however, make it possible to test out these concepts for possible standardization at a  
 34242 later time.

### 34243 **Mutex Attributes and Performance**

34244 Care has been taken to ensure that the default values of the mutex attributes have been defined  
 34245 such that mutexes initialized with the defaults have simple enough semantics so that the locking  
 34246 and unlocking can be done with the equivalent of a test-and-set instruction (plus possibly a few  
 34247 other basic instructions).

34248 There is at least one implementation method that can be used to reduce the cost of testing at  
 34249 lock-time if a mutex has non-default attributes. One such method that an implementation can  
 34250 employ (and this can be made fully transparent to fully conforming POSIX applications) is to  
 34251 secretly pre-lock any mutexes that are initialized to non-default attributes. Any later attempt to  
 34252 lock such a mutex causes the implementation to branch to the “slow path” as if the mutex were  
 34253 unavailable; then, on the slow path, the implementation can do the “real work” to lock a non-  
 34254 default mutex. The underlying unlock operation is more complicated since the implementation  
 34255 never really wants to release the pre-lock on this kind of mutex. This illustrates that, depending  
 34256 on the hardware, there may be certain optimizations that can be used so that whatever mutex  
 34257 attributes are considered “most frequently used” can be processed most efficiently.

### 34258 **Process Shared Memory and Synchronization**

34259 The existence of memory mapping functions in this volume of IEEE Std. 1003.1-200x leads to the  
 34260 possibility that an application may allocate the synchronization objects from this section in  
 34261 memory that is accessed by multiple processes (and therefore, by threads of multiple processes).

34262 In order to permit such usage, while at the same time keeping the usual case (that is, usage  
 34263 within a single process) efficient, a process-shared option has been defined.

34264 If an implementation supports the `_POSIX_THREAD_PROCESS_SHARED` option, then the  
 34265 *process-shared* attribute can be used to indicate that mutexes or condition variables may be  
 34266 accessed by threads of multiple processes.

34267 The default setting of `PTHREAD_PROCESS_PRIVATE` has been chosen for the *process-shared*  
 34268 attribute so that the most efficient forms of these synchronization objects are created by default.

34269 Synchronization variables that are initialized with the `PTHREAD_PROCESS_PRIVATE` *process-*  
 34270 *shared* attribute may only be operated on by threads in the process that initialized them.  
 34271 Synchronization variables that are initialized with the `PTHREAD_PROCESS_SHARED` *process-*  
 34272 *shared* attribute may be operated on by any thread in any process that has access to it. In  
 34273 particular, these processes may exist beyond the lifetime of the initializing process. For example,  
 34274 the following code implements a simple counting semaphore in a mapped file that may be used  
 34275 by many processes.

```
34276 /* sem.h */
34277 struct semaphore {
34278 pthread_mutex_t lock;
34279 pthread_cond_t nonzero;
34280 unsigned count;
34281 };
34282 typedef struct semaphore semaphore_t;
34283 semaphore_t *semaphore_create(char *semaphore_name);
34284 semaphore_t *semaphore_open(char *semaphore_name);
34285 void semaphore_post(semaphore_t *semap);
```

```
34286 void semaphore_wait(semaphore_t *semap);
34287 void semaphore_close(semaphore_t *semap);
34288 /* sem.c */
34289 #include <sys/types.h>
34290 #include <sys/stat.h>
34291 #include <sys/mman.h>
34292 #include <fcntl.h>
34293 #include <pthread.h>
34294 #include "sem.h"
34295 semaphore_t *
34296 semaphore_create(char *semaphore_name)
34297 {
34298 int fd;
34299 semaphore_t *semap;
34300 pthread_mutexattr_t psharedm;
34301 pthread_condattr_t psharedc;
34302 fd = open(semaphore_name, O_RDWR | O_CREAT | O_EXCL, 0666);
34303 if (fd < 0)
34304 return (NULL);
34305 (void) ftruncate(fd, sizeof(semaphore_t));
34306 (void) pthread_mutexattr_init(&psharedm);
34307 (void) pthread_mutexattr_setpshared(&psharedm,
34308 PTHREAD_PROCESS_SHARED);
34309 (void) pthread_condattr_init(&psharedc);
34310 (void) pthread_condattr_setpshared(&psharedc,
34311 PTHREAD_PROCESS_SHARED);
34312 semap = (semaphore_t *) mmap(NULL, sizeof(semaphore_t),
34313 PROT_READ | PROT_WRITE, MAP_SHARED,
34314 fd, 0);
34315 close (fd);
34316 (void) pthread_mutex_init(&semap->lock, &psharedm);
34317 (void) pthread_cond_init(&semap->nonzero, &psharedc);
34318 semap->count = 0;
34319 return (semap);
34320 }
34321 semaphore_t *
34322 semaphore_open(char *semaphore_name)
34323 {
34324 int fd;
34325 semaphore_t *semap;
34326 fd = open(semaphore_name, O_RDWR, 0666);
34327 if (fd < 0)
34328 return (NULL);
34329 semap = (semaphore_t *) mmap(NULL, sizeof(semaphore_t),
34330 PROT_READ | PROT_WRITE, MAP_SHARED,
34331 fd, 0);
34332 close (fd);
34333 return (semap);
34334 }
```

```

34335 void
34336 semaphore_post(semaphore_t *semap)
34337 {
34338 pthread_mutex_lock(&semap->lock);
34339 if (semap->count == 0)
34340 pthread_cond_signal(&semap->nonzero);
34341 semap->count++;
34342 pthread_mutex_unlock(&semap->lock);
34343 }
34344 void
34345 semaphore_wait(semaphore_t *semap)
34346 {
34347 pthread_mutex_lock(&semap->lock);
34348 while (semap->count == 0)
34349 pthread_cond_wait(&semap->nonzero, &semap->lock);
34350 semap->count--;
34351 pthread_mutex_unlock(&semap->lock);
34352 }
34353 void
34354 semaphore_close(semaphore_t *semap)
34355 {
34356 munmap((void *) semap, sizeof(semaphore_t));
34357 }

```

34358 The following code is for three separate processes that create, post, and wait on a semaphore in  
34359 the file **/tmp/semaphore**. Once the file is created, the post and wait programs increment and  
34360 decrement the counting semaphore (waiting and waking as required) even though they did not  
34361 initialize the semaphore.

```

34362 /* create.c */
34363 #include "pthread.h"
34364 #include "sem.h"
34365
34366 int
34367 main()
34368 {
34369 semaphore_t *semap;
34370
34371 semap = semaphore_create("/tmp/semaphore");
34372 if (semap == NULL)
34373 exit(1);
34374 semaphore_close(semap);
34375 return (0);
34376 }
34377
34378 /* post */
34379 #include "pthread.h"
34380 #include "sem.h"
34381
34382 int
34383 main()
34384 {
34385 semaphore_t *semap;

```

```
34382 semap = semaphore_open("/tmp/semaphore");
34383 if (semap == NULL)
34384 exit(1);
34385 semaphore_post(semap);
34386 semaphore_close(semap);
34387 return (0);
34388 }
34389 /* wait */
34390 #include "pthread.h"
34391 #include "sem.h"
34392 int
34393 main()
34394 {
34395 semaphore_t *semap;
34396
34397 semap = semaphore_open("/tmp/semaphore");
34398 if (semap == NULL)
34399 exit(1);
34400 semaphore_wait(semap);
34401 semaphore_close(semap);
34402 return (0);
34403 }
```

**34403 FUTURE DIRECTIONS**

34404 None.

**34405 SEE ALSO**

34406 *pthread\_cond\_destroy()*, *pthread\_create()*, *pthread\_mutex\_destroy()*, *pthread\_mutexattr\_destroy()*, the  
34407 Base Definitions volume of IEEE Std. 1003.1-200x, <**pthread.h**>

**34408 CHANGE HISTORY**

34409 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

**34410 Issue 6**

34411 The *pthread\_mutexattr\_destroy()* and *pthread\_mutexattr\_init()* functions are marked as part of the  
34412 Threads option.

34413 IEEE PASC Interpretation 1003.1c #27 is applied, updating the ERRORS section.

34414 **NAME**

34415 pthread\_mutexattr\_getprioceiling, pthread\_mutexattr\_setprioceiling — get and set prioceiling  
 34416 attribute of mutex attributes object (**REALTIME THREADS**)

34417 **SYNOPSIS**

```
34418 TPP #include <pthread.h>
34419 int pthread_mutexattr_getprioceiling(const pthread_mutexattr_t *restrict attr, |
34420 int *restrict prioceiling); |
34421 int pthread_mutexattr_setprioceiling(pthread_mutexattr_t *attr, |
34422 int prioceiling); |
34423
```

34424 **DESCRIPTION**

34425 The *pthread\_mutexattr\_getprioceiling()* and *pthread\_mutexattr\_setprioceiling()* functions,  
 34426 respectively, get and set the priority ceiling attribute of a mutex attributes object pointed to by  
 34427 *attr* which was previously created by the function *pthread\_mutexattr\_init()*.

34428 The *prioceiling* attribute contains the priority ceiling of initialized mutexes. The values of  
 34429 *prioceiling* are within the maximum range of priorities defined by SCHED\_FIFO.

34430 The *prioceiling* attribute defines the priority ceiling of initialized mutexes, which is the minimum  
 34431 priority level at which the critical section guarded by the mutex is executed. In order to avoid  
 34432 priority inversion, the priority ceiling of the mutex is set to a priority higher than or equal to the  
 34433 highest priority of all the threads that may lock that mutex. The values of *prioceiling* are within  
 34434 the maximum range of priorities defined under the SCHED\_FIFO scheduling policy.

34435 **RETURN VALUE**

34436 Upon successful completion, the *pthread\_mutexattr\_getprioceiling()* and  
 34437 *pthread\_mutexattr\_setprioceiling()* functions shall return zero; otherwise, an error number shall be  
 34438 returned to indicate the error.

34439 **ERRORS**

34440 The *pthread\_mutexattr\_getprioceiling()* and *pthread\_mutexattr\_setprioceiling()* functions may fail if:

34441 [EINVAL] The value specified by *attr* or *prioceiling* is invalid. |

34442 [EPERM] The caller does not have the privilege to perform the operation. |

34443 These functions shall not return an error code of [EINTR]. |

34444 **EXAMPLES**

34445 None.

34446 **APPLICATION USAGE**

34447 None.

34448 **RATIONALE**

34449 None.

34450 **FUTURE DIRECTIONS**

34451 None.

34452 **SEE ALSO**

34453 *pthread\_cond\_destroy()*, *pthread\_create()*, *pthread\_mutex\_destroy()*, the Base Definitions volume of  
 34454 IEEE Std. 1003.1-200x, <pthread.h>

## 34455 CHANGE HISTORY

34456 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34457 Marked as part of the Realtime Threads Feature Group.

## 34458 Issue 6

34459 The `pthread_mutexattr_getprioceiling()` and `pthread_mutexattr_setprioceiling()` functions are  
34460 marked as part of the Thread Priority Protection option.

34461 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
34462 implementation does not support the Thread Priority Protection option.

34463 The [ENOTSUP] error condition has been removed since these functions do not have a *protocol*  
34464 argument.

34465 The **restrict** keyword is added to the `pthread_mutexattr_getprioceiling()` prototype for alignment  
34466 with the ISO/IEC 9899:1999 standard.

34467 **NAME**

34468 pthread\_mutexattr\_getprotocol, pthread\_mutexattr\_setprotocol — get and set protocol attribute  
 34469 of mutex attributes object (**REALTIME THREADS**)

34470 **SYNOPSIS**

```
34471 TPP|TPI #include <pthread.h>
34472 int pthread_mutexattr_getprotocol(const pthread_mutexattr_t *restrict attr, |
34473 int *restrict protocol); |
34474 int pthread_mutexattr_setprotocol(pthread_mutexattr_t *attr, |
34475 int protocol); |
34476
```

34477 **DESCRIPTION**

34478 The *pthread\_mutexattr\_getprotocol()* and *pthread\_mutexattr\_setprotocol()* functions, respectively,  
 34479 get and set the protocol attribute of a mutex attributes object pointed to by *attr* which was  
 34480 previously created by the function *pthread\_mutexattr\_init()*.

34481 The *protocol* attribute defines the protocol to be followed in utilizing mutexes. The value of  
 34482 *protocol* may be one of `PTHREAD_PRIO_NONE`, `PTHREAD_PRIO_INHERIT`, or  
 34483 `PTHREAD_PRIO_PROTECT`, which are defined by the header `<pthread.h>`.

34484 When a thread owns a mutex with the `PTHREAD_PRIO_NONE` *protocol* attribute, its priority  
 34485 and scheduling are not affected by its mutex ownership.

34486 TPI When a thread is blocking higher priority threads because of owning one or more mutexes with  
 34487 the `PTHREAD_PRIO_INHERIT` protocol attribute, it executes at the higher of its priority or the  
 34488 priority of the highest priority thread waiting on any of the mutexes owned by this thread and  
 34489 initialized with this protocol.

34490 TPP When a thread owns one or more mutexes initialized with the `PTHREAD_PRIO_PROTECT`  
 34491 protocol, it executes at the higher of its priority or the highest of the priority ceilings of all the  
 34492 mutexes owned by this thread and initialized with this attribute, regardless of whether other  
 34493 threads are blocked on any of these mutexes or not.

34494 While a thread is holding a mutex which has been initialized with the  
 34495 `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT` protocol attributes, it shall not be  
 34496 subject to being moved to the tail of the scheduling queue at its priority in the event that its  
 34497 original priority is changed, such as by a call to *sched\_setparam()*. Likewise, when a thread  
 34498 unlocks a mutex that has been initialized with the `PTHREAD_PRIO_INHERIT` or  
 34499 `PTHREAD_PRIO_PROTECT` protocol attributes, it shall not be subject to being moved to the tail  
 34500 of the scheduling queue at its priority in the event that its original priority is changed.

34501 If a thread simultaneously owns several mutexes initialized with different protocols, it shall  
 34502 execute at the highest of the priorities that it would have obtained by each of these protocols.

34503 TPI When a thread makes a call to *pthread\_mutex\_lock()*, the mutex was initialized with the protocol  
 34504 attribute having the value `PTHREAD_PRIO_INHERIT`, when the calling thread is blocked  
 34505 because the mutex is owned by another thread, that owner thread shall inherit the priority level  
 34506 of the calling thread as long as it continues to own the mutex. The implementation shall update  
 34507 its execution priority to the maximum of its assigned priority and all its inherited priorities.  
 34508 Furthermore, if this owner thread itself becomes blocked on another mutex, the same priority  
 34509 inheritance effect shall be propagated to this other owner thread, in a recursive manner.

34510 **RETURN VALUE**

34511 Upon successful completion, the *pthread\_mutexattr\_getprotocol()* and  
34512 *pthread\_mutexattr\_setprotocol()* functions shall return zero; otherwise, an error number shall be  
34513 returned to indicate the error.

34514 **ERRORS**

34515 The *pthread\_mutexattr\_setprotocol()* function shall fail if:

34516 [ENOTSUP] The value specified by *protocol* is an unsupported value.

34517 The *pthread\_mutexattr\_getprotocol()* and *pthread\_mutexattr\_setprotocol()* functions may fail if:

34518 [EINVAL] The value specified by *attr* or *protocol* is invalid.

34519 [EPERM] The caller does not have the privilege to perform the operation.

34520 These functions shall not return an error code of [EINTR].

34521 **EXAMPLES**

34522 None.

34523 **APPLICATION USAGE**

34524 None.

34525 **RATIONALE**

34526 None.

34527 **FUTURE DIRECTIONS**

34528 None.

34529 **SEE ALSO**

34530 *pthread\_cond\_destroy()*, *pthread\_create()*, *pthread\_mutex\_destroy()*, the Base Definitions volume of  
34531 IEEE Std. 1003.1-200x, <pthread.h>

34532 **CHANGE HISTORY**

34533 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34534 Marked as part of the Realtime Threads Feature Group.

34535 **Issue 6**

34536 The *pthread\_mutexattr\_getprotocol()* and *pthread\_mutexattr\_setprotocol()* functions are marked as  
34537 part of either the Thread Priority Protection or Threads Priority Inheritance options.

34538 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
34539 implementation does not support the Thread Priority Protection or Threads Priority Inheritance  
34540 options.

34541 The **restrict** keyword is added to the *pthread\_mutexattr\_getprotocol()* prototype for alignment  
34542 with the ISO/IEC 9899:1999 standard.

34543 **NAME**

34544 pthread\_mutexattr\_getpshared, pthread\_mutexattr\_setpshared — get and set process-shared  
 34545 attribute

34546 **SYNOPSIS**

```
34547 THR TSH #include <pthread.h>
34548
34548 int pthread_mutexattr_getpshared(const pthread_mutexattr_t *restrict attr, |
34549 int *restrict pshared);
34550 int pthread_mutexattr_setpshared(pthread_mutexattr_t *attr, |
34551 int pshared);
34552
```

34553 **DESCRIPTION**

34554 The *pthread\_mutexattr\_getpshared()* function obtains the value of the *process-shared* attribute from  
 34555 the attributes object referenced by *attr*. The *pthread\_mutexattr\_setpshared()* function is used to set  
 34556 the *process-shared* attribute in an initialized attributes object referenced by *attr*.

34557 The *process-shared* attribute is set to PTHREAD\_PROCESS\_SHARED to permit a mutex to be  
 34558 operated upon by any thread that has access to the memory where the mutex is allocated, even if  
 34559 the mutex is allocated in memory that is shared by multiple processes. If the *process-shared*  
 34560 attribute is PTHREAD\_PROCESS\_PRIVATE, the mutex shall only be operated upon by threads  
 34561 created within the same process as the thread that initialized the mutex; if threads of differing  
 34562 processes attempt to operate on such a mutex, the behavior is undefined. The default value of  
 34563 the attribute shall be PTHREAD\_PROCESS\_PRIVATE.

34564 **RETURN VALUE**

34565 Upon successful completion, *pthread\_mutexattr\_setpshared()* shall return zero; otherwise, an error  
 34566 number shall be returned to indicate the error.

34567 Upon successful completion, *pthread\_mutexattr\_getpshared()* shall return zero and stores the  
 34568 value of the *process-shared* attribute of *attr* into the object referenced by the *pshared* parameter.  
 34569 Otherwise, an error number shall be returned to indicate the error.

34570 **ERRORS**

34571 The *pthread\_mutexattr\_getpshared()* and *pthread\_mutexattr\_setpshared()* functions may fail if:

34572 [EINVAL] The value specified by *attr* is invalid.

34573 The *pthread\_mutexattr\_setpshared()* function may fail if:

34574 [EINVAL] The new value specified for the attribute is outside the range of legal values  
 34575 for that attribute.

34576 These functions shall not return an error code of [EINTR].

34577 **EXAMPLES**

34578 None.

34579 **APPLICATION USAGE**

34580 None.

34581 **RATIONALE**

34582 None.

34583 **FUTURE DIRECTIONS**

34584 None.

34585 **SEE ALSO**

34586 *pthread\_cond\_destroy()*, *pthread\_create()*, *pthread\_mutex\_destroy()*, *pthread\_mutexattr\_destroy()*, the  
34587 Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

34588 **CHANGE HISTORY**

34589 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

34590 **Issue 6**

34591 The *pthread\_mutexattr\_getpshared()* and *pthread\_mutexattr\_setpshared()* functions are marked as  
34592 part of the Threads and Thread Process-Shared Synchronization options.

34593 The **restrict** keyword is added to the *pthread\_mutexattr\_getpshared()* prototype for alignment  
34594 with the ISO/IEC 9899:1999 standard.

34595 **NAME**

34596 pthread\_mutexattr\_gettype, pthread\_mutexattr\_settype — get or set a mutex type

34597 **SYNOPSIS**

34598 XSI #include &lt;pthread.h&gt;

34599 int pthread\_mutexattr\_gettype(const pthread\_mutexattr\_t \*restrict attr,  
34600 int \*restrict type);

34601 int pthread\_mutexattr\_settype(pthread\_mutexattr\_t \*attr, int type);

34602

34603 **DESCRIPTION**34604 The *pthread\_mutexattr\_gettype()* and *pthread\_mutexattr\_settype()* functions, respectively, get and  
34605 set the mutex *type* attribute. This attribute is set in the *type* parameter to these functions. The  
34606 default value of the *type* attribute is PTHREAD\_MUTEX\_DEFAULT.34607 The type of mutex is contained in the *type* attribute of the mutex attributes. Valid mutex types  
34608 include:

## 34609 PTHREAD\_MUTEX\_NORMAL

34610 This type of mutex does not detect deadlock. A thread attempting to relock this mutex  
34611 without first unlocking it shall deadlock. Attempting to unlock a mutex locked by a  
34612 different thread results in undefined behavior. Attempting to unlock an unlocked mutex  
34613 results in undefined behavior.

## 34614 PTHREAD\_MUTEX\_ERRORCHECK

34615 This type of mutex provides error checking. A thread attempting to relock this mutex  
34616 without first unlocking it shall return with an error. A thread attempting to unlock a mutex  
34617 which another thread has locked shall return with an error. A thread attempting to unlock  
34618 an unlocked mutex shall return with an error.

## 34619 PTHREAD\_MUTEX\_RECURSIVE

34620 A thread attempting to relock this mutex without first unlocking it shall succeed in locking  
34621 the mutex. The relocking deadlock which can occur with mutexes of type  
34622 PTHREAD\_MUTEX\_NORMAL cannot occur with this type of mutex. Multiple locks of this  
34623 mutex require the same number of unlocks to release the mutex before another thread can  
34624 acquire the mutex. A thread attempting to unlock a mutex which another thread has locked  
34625 shall return with an error. A thread attempting to unlock an unlocked mutex shall return  
34626 with an error.

## 34627 PTHREAD\_MUTEX\_DEFAULT

34628 Attempting to recursively lock a mutex of this type results in undefined behavior.  
34629 Attempting to unlock a mutex of this type which was not locked by the calling thread  
34630 results in undefined behavior. Attempting to unlock a mutex of this type which is not  
34631 locked results in undefined behavior. An implementation is allowed to map this mutex to  
34632 one of the other mutex types.34633 **RETURN VALUE**34634 Upon successful completion, the *pthread\_mutexattr\_gettype()* function shall return zero and store  
34635 the value of the *type* attribute of *attr* into the object referenced by the *type* parameter. Otherwise,  
34636 an error shall be returned to indicate the error.34637 If successful, the *pthread\_mutexattr\_settype()* function shall return zero; otherwise, an error  
34638 number shall be returned to indicate the error.

34639 **ERRORS**

34640 The *pthread\_mutexattr\_settype()* function shall fail if:

34641 [EINVAL] The value *type* is invalid.

34642 The *pthread\_mutexattr\_gettype()* and *pthread\_mutexattr\_settype()* functions may fail if:

34643 [EINVAL] The value specified by *attr* is invalid.

34644 These functions shall not return an error code of [EINTR].

34645 **EXAMPLES**

34646 None.

34647 **APPLICATION USAGE**

34648 It is advised that an application should not use a PTHREAD\_MUTEX\_RECURSIVE mutex with  
34649 condition variables because the implicit unlock performed for a *pthread\_cond\_timedwait()* or  
34650 *pthread\_cond\_wait()* may not actually release the mutex (if it had been locked multiple times). If  
34651 this happens, no other thread can satisfy the condition of the predicate.

34652 **RATIONALE**

34653 None.

34654 **FUTURE DIRECTIONS**

34655 None.

34656 **SEE ALSO**

34657 *pthread\_cond\_timedwait()*, *pthread\_cond\_wait()*, the Base Definitions volume of  
34658 IEEE Std. 1003.1-200x, <pthread.h>

34659 **CHANGE HISTORY**

34660 First released in Issue 5.

34661 **Issue 6**

34662 The Open Group corrigenda item U033/3 has been applied. The SYNOPSIS for  
34663 *pthread\_mutexattr\_gettype()* is updated so that the first argument is of type  
34664 **constpthread\_mutexattr\_t\***.

34665 The **restrict** keyword is added to the *pthread\_mutexattr\_gettype()* prototype for alignment with  
34666 the ISO/IEC 9899:1999 standard.

34667 **NAME**

34668 pthread\_mutexattr\_init — initialize mutex attributes object

34669 **SYNOPSIS**

34670 THR #include &lt;pthread.h&gt;

34671 int pthread\_mutexattr\_init(pthread\_mutexattr\_t \*attr);

34672

34673 **DESCRIPTION**34674 Refer to *pthread\_mutexattr\_destroy()*.

34675 **NAME**

34676 pthread\_mutexattr\_setprioceiling — set prioceiling attribute of mutex attributes object  
34677 (**REALTIME THREADS**)

34678 **SYNOPSIS**

34679 TPP #include <pthread.h>

```
34680 int pthread_mutexattr_setprioceiling(pthread_mutexattr_t *attr,
34681 int prioceiling);
```

34682

34683 **DESCRIPTION**

34684 Refer to *pthread\_mutexattr\_getprioceiling()*.

34685 **NAME**

34686 pthread\_mutexattr\_setprotocol — set protocol attribute of mutex attributes object (**REALTIME**  
34687 **THREADS**)

34688 **SYNOPSIS**

```
34689 TPP|TPI #include <pthread.h>
```

```
34690 int pthread_mutexattr_setprotocol(pthread_mutexattr_t *attr,
34691 int protocol);
```

34692

34693 **DESCRIPTION**

34694 Refer to *pthread\_mutexattr\_setprotocol()*.

34695 **NAME**

34696 pthread\_mutexattr\_setpshared — set process-shared attribute

34697 **SYNOPSIS**

34698 THR TSH #include <pthread.h>

```
34699 int pthread_mutexattr_setpshared(pthread_mutexattr_t *attr,
34700 int pshared);
```

34701

34702 **DESCRIPTION**

34703 Refer to *pthread\_mutexattr\_getpshared()*.

34704 **NAME**

34705 pthread\_mutexattr\_settype — set a mutex type

34706 **SYNOPSIS**34707 XSI `#include <pthread.h>`34708 `int pthread_mutexattr_settype(pthread_mutexattr_t *attr, int type);`

34709

34710 **DESCRIPTION**34711 Refer to *pthread\_mutexattr\_gettype()*.

## 34712 NAME

34713 pthread\_once — dynamic package initialization

## 34714 SYNOPSIS

34715 THR #include &lt;pthread.h&gt;

```
34716 int pthread_once(pthread_once_t *once_control,
34717 void (*init_routine)(void));
34718 pthread_once_t once_control = PTHREAD_ONCE_INIT;
34719
```

## 34720 DESCRIPTION

34721 The first call to *pthread\_once()* by any thread in a process, with a given *once\_control*, shall call the  
 34722 *init\_routine* with no arguments. Subsequent calls of *pthread\_once()* with the same *once\_control*  
 34723 shall not call the *init\_routine*. On return from *pthread\_once()*, it is guaranteed that *init\_routine* has  
 34724 completed. The *once\_control* parameter is used to determine whether the associated initialization  
 34725 routine has been called.

34726 The *pthread\_once()* function is not a cancellation point. However, if *init\_routine* is a cancellation  
 34727 point and is canceled, the effect on *once\_control* shall be as if *pthread\_once()* was never called.

34728 The constant PTHREAD\_ONCE\_INIT is defined by the header <pthread.h>.

34729 The behavior of *pthread\_once()* is undefined if *once\_control* has automatic storage duration or is  
 34730 not initialized by PTHREAD\_ONCE\_INIT.

## 34731 RETURN VALUE

34732 Upon successful completion, *pthread\_once()* shall return zero; otherwise, an error number shall  
 34733 be returned to indicate the error.

## 34734 ERRORS

34735 **Notes to Reviewers**34736 *This section with side shading will not appear in the final copy. - Ed.*

34737 D1, XSH, ERN 255 notes that no error is returned for invalid parameters and proposes the  
 34738 following:

34739 [EINVAL] If either *once\_control* or *init\_routine* is invalid.

34740 No errors are defined.

34741 The *pthread\_once()* function shall not return an error code of [EINTR].

## 34742 EXAMPLES

34743 None.

## 34744 APPLICATION USAGE

34745 None.

## 34746 RATIONALE

34747 Some C libraries are designed for dynamic initialization. That is, the global initialization for the  
 34748 library is performed when the first procedure in the library is called. In a single-threaded  
 34749 program, this is normally implemented using a static variable whose value is checked on entry  
 34750 to a routine, as follows:

```
34751 static int random_is_initialized = 0;
34752 extern int initialize_random();
```

```

34753 int random_function()
34754 {
34755 if (random_is_initialized == 0) {
34756 initialize_random();
34757 random_is_initialized = 1;
34758 }
34759 ... /* Operations performed after initialization. */
34760 }

```

34761 To keep the same structure in a multi-threaded program, a new primitive is needed. Otherwise,  
 34762 library initialization has to be accomplished by an explicit call to a library-exported initialization  
 34763 function prior to any use of the library.

34764 For dynamic library initialization in a multi-threaded process, a simple initialization flag is not  
 34765 sufficient; the flag needs to be protected against modification by multiple threads  
 34766 simultaneously calling into the library. Protecting the flag requires the use of a mutex; however,  
 34767 mutexes have to be initialized before they are used. Ensuring that the mutex is only initialized  
 34768 once requires a recursive solution to this problem.

34769 The use of *pthread\_once()* not only supplies an implementation-guaranteed means of dynamic  
 34770 initialization, it provides an aid to the reliable construction of multi-threaded and realtime  
 34771 systems. The preceding example then becomes:

```

34772 #include <pthread.h>
34773 static pthread_once_t random_is_initialized = PTHREAD_ONCE_INIT;
34774 extern int initialize_random();
34775
34776 int random_function()
34777 {
34778 (void) pthread_once(&random_is_initialized, initialize_random);
34779 ... /* Operations performed after initialization. */
34780 }

```

34780 Note that a **pthread\_once\_t** cannot be an array because some compilers do not accept the  
 34781 construct **&<array\_name>**.

#### 34782 FUTURE DIRECTIONS

34783 None.

#### 34784 SEE ALSO

34785 The Base Definitions volume of IEEE Std. 1003.1-200x, **<pthread.h>**

#### 34786 CHANGE HISTORY

34787 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

#### 34788 Issue 6

34789 The *pthread\_once()* function is marked as part of the Threads option.

## 34790 NAME

34791 pthread\_rwlock\_destroy, pthread\_rwlock\_init — destroy and initialize a read-write lock object

## 34792 SYNOPSIS

34793 THR #include <pthread.h>

```
34794 int pthread_rwlock_destroy(pthread_rwlock_t *rwlock);
34795 int pthread_rwlock_init(pthread_rwlock_t *restrict rwlock,
34796 const pthread_rwlockattr_t *restrict attr);
34797
```

## 34798 DESCRIPTION

34799 The *pthread\_rwlock\_destroy()* function shall destroy the read-write lock object referenced by  
 34800 *rwlock* and release any resources used by the lock. The effect of subsequent use of the lock is  
 34801 undefined until the lock is re-initialized by another call to *pthread\_rwlock\_init()*. An  
 34802 implementation may cause *pthread\_rwlock\_destroy()* to set the object referenced by *rwlock* to an  
 34803 invalid value. Results are undefined if *pthread\_rwlock\_destroy()* is called when any thread holds  
 34804 *rwlock*. Attempting to destroy an uninitialized read-write lock results in undefined behavior.

34805 The *pthread\_rwlock\_init()* function shall allocate any resources required to use the read-write  
 34806 lock referenced by *rwlock* and initializes the lock to an unlocked state with attributes referenced  
 34807 by *attr*. If *attr* is NULL, the default read-write lock attributes are used; the effect is the same as  
 34808 passing the address of a default read-write lock attributes object. Once initialized, the lock can be  
 34809 used any number of times without being re-initialized. Results are undefined if  
 34810 *pthread\_rwlock\_init()* is called specifying an already initialized read-write lock. Results are  
 34811 undefined if a read-write lock is used without first being initialized.

34812 If the *pthread\_rwlock\_init()* function fails, *rwlock* is not initialized and the contents of *rwlock* are  
 34813 undefined.

34814 Only the object referenced by *rwlock* may be used for performing synchronization. The result of  
 34815 referring to copies of that object in calls to *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_rdlock()*,  
 34816 *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_tryrdlock()*,  
 34817 *pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_unlock()*, or *pthread\_rwlock\_wrlock()* is undefined.

## 34818 RETURN VALUE

34819 If successful, the *pthread\_rwlock\_destroy()* and *pthread\_rwlock\_init()* functions shall return zero;  
 34820 otherwise, an error number shall be returned to indicate the error.

34821 The [EBUSY] and [EINVAL] error checks, if implemented, act as if they were performed  
 34822 immediately at the beginning of processing for the function and caused an error return prior to  
 34823 modifying the state of the read-write lock specified by *rwlock*.

## 34824 ERRORS

34825 The *pthread\_rwlock\_destroy()* function may fail if:

34826 [EBUSY] The implementation has detected an attempt to destroy the object referenced  
 34827 by *rwlock* while it is locked.

34828 [EINVAL] The value specified by *rwlock* is invalid.

34829 The *pthread\_rwlock\_init()* function shall fail if:

34830 [EAGAIN] The system lacked the necessary resources (other than memory) to initialize  
 34831 another read-write lock.

34832 [ENOMEM] Insufficient memory exists to initialize the read-write lock.

34833 MAN [EPERM] The caller does not have the privilege to perform the operation.

- 34834 The *pthread\_rwlock\_init()* function may fail if:
- 34835 [EBUSY] The implementation has detected an attempt to re-initialize the object  
34836 referenced by *rwlock*, a previously initialized but not yet destroyed read-write  
34837 lock.
- 34838 [EINVAL] The value specified by *attr* is invalid.

### 34839 **Notes to Reviewers**

- 34840 *This section with side shading will not appear in the final copy. - Ed.*
- 34841 D1, XSH, ERN 259 noted that no error return is described for when the argument *rwlock* is  
34842 invalid and proposes adding as a may fail case:
- 34843 [EINVAL] The value specified by *rwlock* is invalid.
- 34844 These functions shall not return an error code of [EINTR].

### 34845 **EXAMPLES**

- 34846 None.

### 34847 **APPLICATION USAGE**

- 34848 None.

### 34849 **RATIONALE**

- 34850 None.

### 34851 **FUTURE DIRECTIONS**

- 34852 None.

### 34853 **SEE ALSO**

- 34854 *pthread\_rwlock\_rdlock()*, *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_timedwrlock()*,  
34855 *pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_unlock()*,  
34856 *pthread\_rwlock\_wrlock()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**pthread.h**>

### 34857 **CHANGE HISTORY**

- 34858 First released in Issue 5.

### 34859 **Issue 6**

- 34860 The following changes are made for alignment with IEEE Std. 1003.1j-2000:
- 34861 • The margin code in the SYNOPSIS s changed to RWL and the initializer macro is deleted.
- 34862 • The DESCRIPTION is updated as follows:
- 34863 — It explicitly notes allocation of resources upon initialization of a read-write lock object.
- 34864 — A paragraph is added specifying that copies of read-write lock objects may not be used.
- 34865 • An [EINVAL] error is added to the ERRORS section for *pthread\_rwlock\_init()*, indicating that  
34866 the *rwlock* value is invalid.
- 34867 • The SEE ALSO section is updated.
- 34868 The **restrict** keyword is added to the *pthread\_rwlock\_init()* prototype for alignment with the  
34869 ISO/IEC 9899:1999 standard.

34870 **NAME**

34871 pthread\_rwlock\_init — initialize a read-write lock object

34872 **SYNOPSIS**

34873 THR #include <pthread.h>

```
34874 int pthread_rwlock_init(pthread_rwlock_t *restrict rwlock,
34875 const pthread_rwlockattr_t *restrict attr);
34876
```

34877 **DESCRIPTION**

34878 Refer to *pthread\_rwlock\_destroy()*.

34879 **NAME**

34880 pthread\_rwlock\_rdlock, pthread\_rwlock\_tryrdlock — lock a read-write lock object for reading

34881 **SYNOPSIS**

34882 THR #include &lt;pthread.h&gt;

34883 int pthread\_rwlock\_rdlock(pthread\_rwlock\_t \*rwlock);

34884 int pthread\_rwlock\_tryrdlock(pthread\_rwlock\_t \*rwlock);

34885

34886 **DESCRIPTION**

34887 The *pthread\_rwlock\_rdlock()* function shall apply a read lock to the read-write lock referenced by  
 34888 *rwlock*. The calling thread acquires the read lock if a writer does not hold the lock and there are  
 34889 TPS no writers blocked on the lock. If the Thread Execution Scheduling option is supported, and the  
 34890 threads involved in the lock are executing with the scheduling policies SCHED\_FIFO,  
 34891 SCHED\_RR, or SCHED\_SPORADIC, the calling thread shall not acquire the lock if a writer  
 34892 holds the lock or if writers of higher or equal priority are blocked on the lock; otherwise, the  
 34893 calling thread shall acquire the lock. If the Thread Execution Scheduling option is not supported,  
 34894 it is implementation-defined whether the calling thread acquires the lock when a writer does not  
 34895 hold the lock and there are writers blocked on the lock. If a writer holds the lock, the calling  
 34896 thread shall not acquire the read lock. If the read lock is not acquired, the calling thread blocks  
 34897 (that is, it does not return from the *pthread\_rwlock\_rdlock()* call) until it can acquire the lock. The  
 34898 calling thread may deadlock if at the time the call is made it holds a write lock.

34899 A thread may hold multiple concurrent read locks on *rwlock* (that is, successfully call the  
 34900 *pthread\_rwlock\_rdlock()* function *n* times). If so, the application shall ensure that the thread  
 34901 performs matching unlocks (that is, it calls the *pthread\_rwlock\_unlock()* function *n* times).

34902 The maximum number of simultaneous read locks that an implementation guarantees can be  
 34903 applied to a read-write lock shall be implementation-defined. The *pthread\_rwlock\_rdlock()*  
 34904 function may fail if this maximum would be exceeded.

34905 The *pthread\_rwlock\_tryrdlock()* function shall apply a read lock as in the *pthread\_rwlock\_rdlock()*  
 34906 function, with the exception that the function shall fail if the equivalent *pthread\_rwlock\_rdlock()*  
 34907 call would have blocked the calling thread. In no case does the *pthread\_rwlock\_tryrdlock()*  
 34908 function ever block; it always either acquires the lock or fails and returns immediately.

34909 Results are undefined if any of these functions are called with an uninitialized read-write lock.

34910 If a signal is delivered to a thread waiting for a read-write lock for reading, upon return from the  
 34911 signal handler the thread resumes waiting for the read-write lock for reading as if it was not  
 34912 interrupted.

34913 **RETURN VALUE**

34914 If successful, the *pthread\_rwlock\_rdlock()* function shall return zero; otherwise, an error number  
 34915 shall be returned to indicate the error.

34916 The *pthread\_rwlock\_tryrdlock()* function shall return zero if the lock for reading on the read-write  
 34917 lock object referenced by *rwlock* is acquired. Otherwise, an error number shall be returned to  
 34918 indicate the error.

34919 **ERRORS**

34920 The *pthread\_rwlock\_tryrdlock()* function shall fail if:

34921 [EBUSY] The read-write lock could not be acquired for reading because a writer holds  
 34922 the lock or a writer with the appropriate priority was blocked on it.

34923 The *pthread\_rwlock\_rdlock()* and *pthread\_rwlock\_tryrdlock()* functions may fail if:

34924 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
34925 object.

34926 [EAGAIN] The read lock could not be acquired because the maximum number of read  
34927 locks for *rwlock* has been exceeded.

34928 The *pthread\_rwlock\_rdlock()* function may fail if:

34929 [EDEADLK] The current thread already owns the read-write lock for writing.

34930 These functions shall not return an error code of [EINTR].

#### 34931 EXAMPLES

34932 None.

#### 34933 APPLICATION USAGE

34934 Applications using these functions may be subject to priority inversion, as discussed in the Base  
34935 Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

#### 34936 RATIONALE

34937 None.

#### 34938 FUTURE DIRECTIONS

34939 None.

#### 34940 SEE ALSO

34941 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_timedrdlock()*,  
34942 *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_unlock()*,  
34943 *pthread\_rwlock\_wrlock()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

#### 34944 CHANGE HISTORY

34945 First released in Issue 5.

#### 34946 Issue 6

34947 The following changes are made for alignment with IEEE Std. 1003.1j-2000:

- 34948 • The margin code in the SYNOPSIS is changed to RWL.
- 34949 • The DESCRIPTION is updated as follows:
  - 34950 — Conditions under which writers have precedence over readers are specified.
  - 34951 — Failure of *pthread\_rwlock\_tryrdlock()* is clarified.
  - 34952 — A paragraph on the maximum number of read locks is added.
- 34953 • In the ERRORS sections, [EBUSY] is modified to take into account write priority, and  
34954 [EDEADLK] is deleted as a *pthread\_rwlock\_tryrdlock()* error.
- 34955 • The SEE ALSO section is updated.

34956 **NAME**

34957 pthread\_rwlock\_timedrdlock — lock a read-write lock for reading

34958 **SYNOPSIS**

34959 THR TMO #include &lt;pthread.h&gt;

34960 #include &lt;time.h&gt;

34961 int pthread\_rwlock\_timedrdlock(pthread\_rwlock\_t \*restrict *rwlock*,34962 const struct timespec \*restrict *abs\_timeout*);

34963

34964 **DESCRIPTION**

34965 The *pthread\_rwlock\_timedrdlock()* function applies a read lock to the read-write lock referenced  
 34966 by *rwlock* as in the *pthread\_rwlock\_rdlock()* function. However, if the lock cannot be acquired  
 34967 without waiting for other threads to unlock the lock, this wait shall be terminated when the  
 34968 specified timeout expires. The timeout expires when the absolute time specified by *abs\_timeout*  
 34969 passes, as measured by the clock on which timeouts are based (that is, when the value of that  
 34970 clock equals or exceeds *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already  
 34971 been passed at the time of the call.

34972 TMR If the Timers option is supported, the timeout is based on the `CLOCK_REALTIME` clock; if the  
 34973 Timers option is not supported, the timeout is based on the system clock as returned by the  
 34974 *time()* function. The resolution of the timeout is the resolution of the clock on which it is based.  
 34975 The `timespec` data type is defined as a structure in the `<time.h>` header. Under no circumstances  
 34976 shall the function fail with a timeout if the lock can be acquired immediately. The validity of the  
 34977 *abs\_timeout* parameter need not be checked if the lock can be immediately acquired.

34978 If a signal that causes a signal handler to be executed is delivered to a thread blocked on a read-  
 34979 write lock via a call to *pthread\_rwlock\_timedrdlock()*, upon return from the signal handler the  
 34980 thread shall resume waiting for the lock as if it was not interrupted.

34981 The calling thread may deadlock if at the time the call is made it holds a write lock on *rwlock*.  
 34982 The results are undefined if this function is called with an uninitialized read-write lock.

34983 **RETURN VALUE**

34984 The *pthread\_rwlock\_timedrdlock()* function shall return zero if the lock for reading on the read-  
 34985 write lock object referenced by *rwlock* is acquired. Otherwise, an error number shall be returned  
 34986 to indicate the error.

34987 **ERRORS**34988 The *pthread\_rwlock\_timedrdlock()* function shall fail if:

34989 [ETIMEDOUT] The lock could not be acquired before the specified timeout expired.

34990 The *pthread\_rwlock\_timedrdlock()* function may fail if:34991 [EAGAIN] The read lock could not be acquired because the maximum number of read  
34992 locks for lock would be exceeded.34993 [EDEADLK] The calling thread already holds a write lock on *rwlock*.34994 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
34995 object, or the *abs\_timeout* nanosecond value is less than zero or greater than or  
34996 equal to 1 000 million.

34997 This function shall not return an error code of [EINTR].

## 34998 EXAMPLES

34999 None.

## 35000 APPLICATION USAGE

35001 Applications using this function may be subject to priority inversion, as discussed in the Base  
35002 Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

35003 The *pthread\_rwlock\_timedrdlock()* function is part of the Threads and Timeouts options and need  
35004 not be provided on all implementations.

## 35005 RATIONALE

35006 None.

## 35007 FUTURE DIRECTIONS

35008 None.

## 35009 SEE ALSO

35010 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
35011 *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*,  
35012 *pthread\_rwlock\_unlock()*, *pthread\_rwlock\_wrlock()*, the Base Definitions volume of  
35013 IEEE Std. 1003.1-200x, <pthread.h>, <time.h>

## 35014 CHANGE HISTORY

35015 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

35016 **NAME**

35017 pthread\_rwlock\_timedwrlock — lock a read-write lock for writing

35018 **SYNOPSIS**

35019 THR TMO #include &lt;pthread.h&gt;

35020 #include &lt;time.h&gt;

35021 int pthread\_rwlock\_timedwrlock(pthread\_rwlock\_t \*restrict *rwlock*,35022 const struct timespec \*restrict *abs\_timeout*);

35023

35024 **DESCRIPTION**

35025 The *pthread\_rwlock\_timedwrlock()* function applies a write lock to the read-write lock referenced  
 35026 by *rwlock* as in the *pthread\_rwlock\_wrlock()* function. However, if the lock cannot be acquired  
 35027 without waiting for other threads to unlock the lock, this wait shall be terminated when the  
 35028 specified timeout expires. The timeout expires when the absolute time specified by *abs\_timeout*  
 35029 passes, as measured by the clock on which timeouts are based (that is, when the value of that  
 35030 clock equals or exceeds *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already  
 35031 been passed at the time of the call.

35032 TMR If the Timers option is supported, the timeout is based on the `CLOCK_REALTIME` clock; if the  
 35033 Timers option is not supported, the timeout is based on the system clock as returned by the  
 35034 *time()* function. The resolution of the timeout is the resolution of the clock on which it is based.  
 35035 The `timespec` data type is defined as a structure in the `<time.h>` header. Under no circumstances  
 35036 shall the function fail with a timeout if the lock can be acquired immediately. The validity of the  
 35037 *abs\_timeout* parameter need not be checked if the lock can be immediately acquired.

35038 If a signal that causes a signal handler to be executed is delivered to a thread blocked on a read-  
 35039 write lock via a call to *pthread\_rwlock\_timedwrlock()*, upon return from the signal handler the  
 35040 thread shall resume waiting for the lock as if it was not interrupted.

35041 The calling thread may deadlock if at the time the call is made it holds the read-write lock. The  
 35042 results are undefined if this function is called with an uninitialized read-write lock.

35043 **RETURN VALUE**

35044 The *pthread\_rwlock\_timedwrlock()* function shall return zero if the lock for writing on the read-  
 35045 write lock object referenced by *rwlock* is acquired. Otherwise, an error number shall be returned  
 35046 to indicate the error.

35047 **ERRORS**35048 The *pthread\_rwlock\_timedwrlock()* function shall fail if:

35049 [ETIMEDOUT] The lock could not be acquired before the specified timeout expired.

35050 The *pthread\_rwlock\_timedwrlock()* function may fail if:35051 [EDEADLK] The calling thread already holds the *rwlock*.

35052 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
 35053 object, or the *abs\_timeout* nanosecond value is less than zero or greater than or  
 35054 equal to 1 000 million.

35055 This function shall not return an error code of [EINTR].

35056 **EXAMPLES**

35057           None.

35058 **APPLICATION USAGE**

35059           Applications using this function may be subject to priority inversion, as discussed in the Base  
35060           Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

35061           The *pthread\_rwlock\_timedwrlock()* function is part of the Threads and Timeouts options and need  
35062           not be provided on all implementations.

35063 **RATIONALE**

35064           None.

35065 **FUTURE DIRECTIONS**

35066           None.

35067 **SEE ALSO**

35068           *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
35069           *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*,  
35070           *pthread\_rwlock\_unlock()*, *pthread\_rwlock\_wrlock()*, the Base Definitions volume of  
35071           IEEE Std. 1003.1-200x, <pthread.h>, <time.h>

35072 **CHANGE HISTORY**

35073           First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

35074 **NAME**

35075 pthread\_rwlock\_tryrdlock — lock a read-write lock object for reading

35076 **SYNOPSIS**

35077 THR #include &lt;pthread.h&gt;

35078 int pthread\_rwlock\_tryrdlock(pthread\_rwlock\_t \*rwlock);

35079

35080 **DESCRIPTION**35081 Refer to *pthread\_rwlock\_rdlock()*.

## 35082 NAME

35083 pthread\_rwlock\_trywrlock, pthread\_rwlock\_wrlock — lock a read-write lock object for writing

## 35084 SYNOPSIS

```
35085 THR #include <pthread.h>
```

```
35086 int pthread_rwlock_trywrlock(pthread_rwlock_t *rwlock);
```

```
35087 int pthread_rwlock_wrlock(pthread_rwlock_t *rwlock);
```

35088

## 35089 DESCRIPTION

35090 The *pthread\_rwlock\_trywrlock()* function shall apply a write lock like the *pthread\_rwlock\_wrlock()*  
35091 function, with the exception that the function shall fail if any thread currently holds *rwlock* (for  
35092 reading or writing).

35093 The *pthread\_rwlock\_wrlock()* function shall apply a write lock to the read-write lock referenced  
35094 by *rwlock*. The calling thread acquires the write lock if no other thread (reader or writer) holds  
35095 the read-write lock *rwlock*. Otherwise, the thread blocks (that is, does not return from the  
35096 *pthread\_rwlock\_wrlock()* call) until it can acquire the lock. The calling thread may deadlock if at  
35097 the time the call is made it holds the read-write lock (whether a read or write lock).

35098 Implementations are allowed to favor writers over readers to avoid writer starvation.

35099 Results are undefined if any of these functions are called with an uninitialized read-write lock.

35100 If a signal is delivered to a thread waiting for a read-write lock for writing, upon return from the  
35101 signal handler the thread resumes waiting for the read-write lock for writing as if it was not  
35102 interrupted.

## 35103 RETURN VALUE

35104 The *pthread\_rwlock\_trywrlock()* function shall return zero if the lock for writing on the read-write  
35105 lock object referenced by *rwlock* is acquired. Otherwise, an error number shall be returned to  
35106 indicate the error.

35107 If successful, the *pthread\_rwlock\_wrlock()* function shall return zero; otherwise, an error number  
35108 shall be returned to indicate the error.

## 35109 ERRORS

35110 The *pthread\_rwlock\_trywrlock()* function shall fail if:

35111 [EBUSY] The read-write lock could not be acquired for writing because it was already  
35112 locked for reading or writing.

35113 The *pthread\_rwlock\_trywrlock()* and *pthread\_rwlock\_wrlock()* functions may fail if:

35114 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
35115 object.

35116 The *pthread\_rwlock\_wrlock()* function may fail if:

35117 [EDEADLK] The current thread already owns the read-write lock for writing or reading.

35118 These functions shall not return an error code of [EINTR].

35119 **EXAMPLES**

35120 None.

35121 **APPLICATION USAGE**

35122 Applications using these functions may be subject to priority inversion, as discussed in the Base  
35123 Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

35124 **RATIONALE**

35125 None.

35126 **FUTURE DIRECTIONS**

35127 None.

35128 **SEE ALSO**

35129 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
35130 *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_tryrdlock()*,  
35131 *pthread\_rwlock\_unlock()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**pthread.h**>

35132 **CHANGE HISTORY**

35133 First released in Issue 5.

35134 **Issue 6**

35135 The following changes are made for alignment with IEEE Std. 1003.1j-2000:

- 35136
- The margin code in the SYNOPSIS is changed to RWL.
  - The [EDEADLK] error is deleted as a *pthread\_rwlock\_trywrlock()* error.
  - The SEE ALSO section is updated.
- 35137
- 35138

35139 **NAME**

35140 pthread\_rwlock\_unlock — unlock a read-write lock object

35141 **SYNOPSIS**

35142 THR #include &lt;pthread.h&gt;

35143 int pthread\_rwlock\_unlock(pthread\_rwlock\_t \*rwlock);

35144

35145 **DESCRIPTION**

35146 The *pthread\_rwlock\_unlock()* function is called to release a lock held on the read-write lock object  
35147 referenced by *rwlock*. Results are undefined if the read-write lock *rwlock* is not held by the  
35148 calling thread.

35149 If this function is called to release a read lock from the read-write lock object and there are other  
35150 read locks currently held on this read-write lock object, the read-write lock object remains in the  
35151 read locked state. If this function releases the last read lock for this read-write lock object, the  
35152 read-write lock object shall be put in the unlocked state with no owners.

35153 If this function is called to release a write lock for this read-write lock object, the read-write lock  
35154 object shall be put in the unlocked state.

35155 If there are threads blocked on the lock when it becomes available, the scheduling policy is used  
35156 to determine which thread(s) shall acquire the lock. If the Thread Execution Scheduling option is  
35157 supported, when threads executing with the scheduling policies SCHED\_FIFO, SCHED\_RR, or  
35158 SCHED\_SPORADIC are waiting on the lock, they will acquire the lock in priority order when  
35159 the lock becomes available. For equal priority threads, write locks take precedence over read  
35160 locks. If the Thread Execution Scheduling option is not supported, it is implementation-defined  
35161 whether write locks take precedence over read locks.

35162 Results are undefined if any of these functions are called with an uninitialized read-write lock.

35163 **RETURN VALUE**

35164 If successful, the *pthread\_rwlock\_unlock()* function shall return zero; otherwise, an error number  
35165 shall be returned to indicate the error.

35166 **ERRORS**

35167 The *pthread\_rwlock\_unlock()* function may fail if:

35168 [EINVAL] The value specified by *rwlock* does not refer to an initialized read-write lock  
35169 object.

35170 [EPERM] The current thread does not hold a lock on the read-write lock.

35171 The *pthread\_rwlock\_unlock()* function shall not return an error code of [EINTR].

35172 **EXAMPLES**

35173 None.

35174 **APPLICATION USAGE**

35175 None.

35176 **RATIONALE**

35177 None.

35178 **FUTURE DIRECTIONS**

35179 None.

35180 **SEE ALSO**

35181 *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
35182 *pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_tryrdlock()*,  
35183 *pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_wrlock()*, the Base Definitions volume of  
35184 IEEE Std. 1003.1-200x, <pthread.h>

35185 **CHANGE HISTORY**

35186 First released in Issue 5.

35187 **Issue 6**

35188 The following changes are made for alignment with IEEE Std. 1003.1j-2000:

- 35189 • The margin code in the SYNOPSIS is changed to RWL.
- 35190 • The DESCRIPTION is updated as follows:
  - 35191 — The conditions under which writers have precedence over readers are specified.
  - 35192 — The concept of read-write lock owner is deleted.
- 35193 • The SEE ALSO section is updated.

35194 **NAME**

35195 pthread\_rwlock\_wrlock — lock a read-write lock object for writing

35196 **SYNOPSIS**

35197 THR #include <pthread.h>

35198 int pthread\_rwlock\_wrlock(pthread\_rwlock\_t \**rwlock*);

35199

35200 **DESCRIPTION**

35201 Refer to *pthread\_rwlock\_trywrlock()*.

35202 **NAME**

35203 pthread\_rwlockattr\_destroy, pthread\_rwlockattr\_init — destroy and initialize read-write lock  
 35204 attributes object

35205 **SYNOPSIS**

```
35206 THR #include <pthread.h>
35207
35207 int pthread_rwlockattr_destroy(pthread_rwlockattr_t *attr);
35208 int pthread_rwlockattr_init(pthread_rwlockattr_t *attr);
35209
```

35210 **DESCRIPTION**

35211 The *pthread\_rwlockattr\_destroy()* function shall destroy a read-write lock attributes object. The  
 35212 effect of subsequent use of the object is undefined until the object is re-initialized by another call  
 35213 to *pthread\_rwlockattr\_init()*. An implementation may cause *pthread\_rwlockattr\_destroy()* to set  
 35214 the object referenced by *attr* to an invalid value.

35215 The *pthread\_rwlockattr\_init()* function shall initialize a read-write lock attributes object *attr* with  
 35216 the default value for all of the attributes defined by the implementation.

35217 Results are undefined if *pthread\_rwlockattr\_init()* is called specifying an already initialized read-  
 35218 write lock attributes object.

35219 After a read-write lock attributes object has been used to initialize one or more read-write locks,  
 35220 any function affecting the attributes object (including destruction) does not affect any previously  
 35221 initialized read-write locks.

35222 **RETURN VALUE**

35223 If successful, the *pthread\_rwlockattr\_destroy()* and *pthread\_rwlockattr\_init()* functions shall return  
 35224 zero; otherwise, an error number shall be returned to indicate the error.

35225 **ERRORS**

35226 The *pthread\_rwlockattr\_destroy()* function may fail if:

35227 [EINVAL] The value specified by *attr* is invalid.

35228 The *pthread\_rwlockattr\_init()* function shall fail if:

35229 [ENOMEM] Insufficient memory exists to initialize the read-write lock attributes object.

35230 These functions shall not return an error code of [EINTR].

35231 **EXAMPLES**

35232 None.

35233 **APPLICATION USAGE**

35234 None.

35235 **RATIONALE**

35236 None.

35237 **FUTURE DIRECTIONS**

35238 None.

35239 **SEE ALSO**

35240 *pthread\_rwlock\_init()*, *pthread\_rwlockattr\_getpshared()*, *pthread\_rwlockattr\_setpshared()*, the Base  
 35241 Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

35242 **CHANGE HISTORY**

35243 First released in Issue 5.

35244 **Issue 6**

35245 The following changes are made for alignment with IEEE Std. 1003.1j-2000:

- 35246 • The margin code in the SYNOPSIS is changed to RWL.
- 35247 • The SEE ALSO section is updated.

35248 **NAME**

35249 pthread\_rwlockattr\_getpshared, pthread\_rwlockattr\_setpshared — get and set process-shared  
 35250 attribute of read-write lock attributes object

35251 **SYNOPSIS**

```
35252 THR TSH #include <pthread.h>
35253
35253 int pthread_rwlockattr_getpshared(const pthread_rwlockattr_t *restrict attr,
35254 int *restrict pshared);
35255 int pthread_rwlockattr_setpshared(pthread_rwlockattr_t *attr,
35256 int pshared);
35257
```

35258 **DESCRIPTION**

35259 The *process-shared* attribute is set to PTHREAD\_PROCESS\_SHARED to permit a read-write lock  
 35260 to be operated upon by any thread that has access to the memory where the read-write lock is  
 35261 allocated, even if the read-write lock is allocated in memory that is shared by multiple processes.  
 35262 If the *process-shared* attribute is PTHREAD\_PROCESS\_PRIVATE, the read-write lock shall only  
 35263 be operated upon by threads created within the same process as the thread that initialized the  
 35264 read-write lock; if threads of differing processes attempt to operate on such a read-write lock,  
 35265 the behavior is undefined. The default value of the *process-shared* attribute is  
 35266 PTHREAD\_PROCESS\_PRIVATE.

35267 The *pthread\_rwlockattr\_getpshared()* function obtains the value of the *process-shared* attribute from  
 35268 the initialized attributes object referenced by *attr*. The *pthread\_rwlockattr\_setpshared()* function is  
 35269 used to set the *process-shared* attribute in an initialized attributes object referenced by *attr*.

35270 Additional attributes, their default values, and the names of the associated functions to get and  
 35271 set those attribute values are implementation-defined.

35272 **RETURN VALUE**

35273 Upon successful completion, the *pthread\_rwlockattr\_getpshared()* shall return zero and store the  
 35274 value of the *process-shared* attribute of *attr* into the object referenced by the *pshared* parameter.  
 35275 Otherwise, an error number shall be returned to indicate the error.

35276 If successful, the *pthread\_rwlockattr\_setpshared()* function shall return zero; otherwise, an error  
 35277 number shall be returned to indicate the error.

35278 **ERRORS**

35279 The *pthread\_rwlockattr\_getpshared()* and *pthread\_rwlockattr\_setpshared()* functions may fail if:

35280 [EINVAL] The value specified by *attr* is invalid.

35281 The *pthread\_rwlockattr\_setpshared()* function may fail if:

35282 [EINVAL] The new value specified for the attribute is outside the range of legal values  
 35283 for that attribute.

35284 These functions shall not return an error code of [EINTR].

35285 **EXAMPLES**

35286 None.

35287 **APPLICATION USAGE**

35288 None.

35289 **RATIONALE**

35290 None.

35291 **FUTURE DIRECTIONS**

35292 None.

35293 **SEE ALSO**

35294 *pthread\_rwlock\_init()*, *pthread\_rwlockattr\_destroy()*, *pthread\_rwlockattr\_init()*, the Base Definitions  
35295 volume of IEEE Std. 1003.1-200x, <pthread.h>

35296 **CHANGE HISTORY**

35297 First released in Issue 5.

35298 **Issue 6**

35299 The following changes are made for alignment with IEEE Std. 1003.1j-2000:

- 35300 • The margin code in the SYNOPSIS is changed to RWL TSH.
- 35301 • The DESCRIPTION notes that additional attributes are implementation-defined.
- 35302 • The SEE ALSO section is updated.

35303 The **restrict** keyword is added to the *pthread\_rwlockattr\_getpshared()* prototype for alignment  
35304 with the ISO/IEC 9899:1999 standard.

35305 **NAME**

35306 pthread\_rwlockattr\_init — initialize read-write lock attributes object

35307 **SYNOPSIS**

35308 XSI #include &lt;pthread.h&gt;

35309 int pthread\_rwlockattr\_init(pthread\_rwlockattr\_t \*attr);

35310

35311 **DESCRIPTION**35312 Refer to *pthread\_rwlockattr\_destroy()*.

35313 **NAME**

35314 pthread\_rwlockattr\_setpshared — set process-shared attribute of read-write lock attributes  
35315 object

35316 **SYNOPSIS**

35317 XSI #include <pthread.h>

```
35318 int pthread_rwlockattr_setpshared(pthread_rwlockattr_t *attr,
35319 int pshared);
```

35320

35321 **DESCRIPTION**

35322 Refer to *pthread\_rwlockattr\_getpshared()*.

35323 **NAME**

35324 pthread\_self — get calling thread's ID

35325 **SYNOPSIS**

35326 THR #include &lt;pthread.h&gt;

35327 pthread\_t pthread\_self(void);

35328

35329 **DESCRIPTION**35330 The *pthread\_self()* function shall return the thread ID of the calling thread.35331 **RETURN VALUE**

35332 Refer to the DESCRIPTION.

35333 **ERRORS**

35334 No errors are defined.

35335 The *pthread\_self()* function shall not return an error code of [EINTR].35336 **EXAMPLES**

35337 None.

35338 **APPLICATION USAGE**

35339 None.

35340 **RATIONALE**35341 The *pthread\_self()* function provides a capability similar to the *getpid()* function for processes  
35342 and the rationale is the same: the creation call does not provide the thread ID to the created  
35343 thread.35344 **FUTURE DIRECTIONS**

35345 None.

35346 **SEE ALSO**35347 *pthread\_create()*, *pthread\_equal()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
35348 <pthread.h>35349 **CHANGE HISTORY**

35350 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

35351 **Issue 6**35352 The *pthread\_self()* function is marked as part of the Threads option.

## 35353 NAME

35354 pthread\_setcancelstate, pthread\_setcanceltype, pthread\_testcancel — set cancelability state

## 35355 SYNOPSIS

35356 THR #include &lt;pthread.h&gt;

35357 int pthread\_setcancelstate(int *state*, int \**oldstate*);35358 int pthread\_setcanceltype(int *type*, int \**oldtype*);

35359 void pthread\_testcancel(void);

35360

## 35361 DESCRIPTION

35362 The *pthread\_setcancelstate()* function shall atomically both set the calling thread's cancelability  
 35363 state to the indicated *state* and return the previous cancelability state at the location referenced  
 35364 by *oldstate*. Legal values for *state* are PTHREAD\_CANCEL\_ENABLE and  
 35365 PTHREAD\_CANCEL\_DISABLE.

35366 The *pthread\_setcanceltype()* function shall atomically both set the calling thread's cancelability  
 35367 type to the indicated *type* and return the previous cancelability type at the location referenced by  
 35368 *oldtype*. Legal values for *type* are PTHREAD\_CANCEL\_DEFERRED and  
 35369 PTHREAD\_CANCEL\_ASYNCHRONOUS.

35370 The cancelability state and type of any newly created threads, including the thread in which  
 35371 *main()* was first invoked, shall be PTHREAD\_CANCEL\_ENABLE and  
 35372 PTHREAD\_CANCEL\_DEFERRED respectively.

35373 The *pthread\_testcancel()* function shall create a cancelation point in the calling thread. The  
 35374 *pthread\_testcancel()* function shall have no effect if cancelability is disabled.

## 35375 RETURN VALUE

35376 If successful, the *pthread\_setcancelstate()* and *pthread\_setcanceltype()* functions shall return zero;  
 35377 otherwise, an error number shall be returned to indicate the error.

## 35378 ERRORS

35379 The *pthread\_setcancelstate()* function may fail if:

35380 [EINVAL] The specified *state* is not PTHREAD\_CANCEL\_ENABLE or  
 35381 PTHREAD\_CANCEL\_DISABLE.

35382 The *pthread\_setcanceltype()* function may fail if:

35383 [EINVAL] The specified *type* is not PTHREAD\_CANCEL\_DEFERRED or  
 35384 PTHREAD\_CANCEL\_ASYNCHRONOUS.

35385 These functions shall not return an error code of [EINTR].

## 35386 EXAMPLES

35387 None.

## 35388 APPLICATION USAGE

35389 None.

## 35390 RATIONALE

35391 The *pthread\_setcancelstate()* and *pthread\_setcanceltype()* functions are used to control the points at  
 35392 which a thread may be asynchronously canceled. For cancelation control to be usable in modular  
 35393 fashion, some rules need to be followed.

35394 An object can be considered to be a generalization of a procedure. It is a set of procedures and  
 35395 global variables written as a unit and called by clients not known by the object. Objects may  
 35396 depend on other objects.

35397 First, cancelability should only be disabled on entry to an object, never explicitly enabled. On  
35398 exit from an object, the cancelability state should always be restored to its value on entry to the  
35399 object.

35400 This follows from a modularity argument: if the client of an object (or the client of an object that  
35401 uses that object) has disabled cancelability, it is because the client does not want to be concerned  
35402 about cleaning up if the thread is canceled while executing some sequence of actions. If an object  
35403 is called in such a state and it enables cancelability and a cancelation request is pending for that  
35404 thread, then the thread is canceled, contrary to the wish of the client that disabled.

35405 Second, the cancelability type may be explicitly set to either *deferred* or *asynchronous* upon entry  
35406 to an object. But as with the cancelability state, on exit from an object the cancelability type  
35407 should always be restored to its value on entry to the object.

35408 Finally, only functions that are cancel-safe may be called from a thread that is asynchronously  
35409 cancelable.

#### 35410 **FUTURE DIRECTIONS**

35411 None.

#### 35412 **SEE ALSO**

35413 *pthread\_cancel()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

#### 35414 **CHANGE HISTORY**

35415 First released in Issue 5. Included for alignment with the POSIX Threads Extension.

#### 35416 **Issue 6**

35417 The *pthread\_setcancelstate()*, *pthread\_setcanceltype()*, and *pthread\_testcancel()* functions are marked  
35418 as part of the Threads option.

35419 **NAME**

35420 pthread\_setconcurrency — set level of concurrency

35421 **SYNOPSIS**

35422 XSI #include <pthread.h>

35423 int pthread\_setconcurrency(int new\_level);

35424

35425 **DESCRIPTION**

35426 Refer to *pthread\_getconcurrency()*.

35427 **NAME**

35428 pthread\_setschedparam — dynamic thread scheduling parameters access (**REALTIME**  
35429 **THREADS**)

35430 **SYNOPSIS**

35431 TPS #include <pthread.h>

35432 int pthread\_setschedparam(pthread\_t *thread*, int *policy*,  
35433 const struct sched\_param \**param*);

35434

35435 **DESCRIPTION**

35436 Refer to *pthread\_getschedparam()*.

35437 **NAME**

35438 pthread\_setspecific — thread-specific data management

35439 **SYNOPSIS**

35440 THR #include <pthread.h>

35441 int pthread\_setspecific(pthread\_key\_t key, const void \*value);

35442

35443 **DESCRIPTION**

35444 Refer to *pthread\_getspecific()*.

35445 **NAME**

35446 pthread\_sigmask, sigprocmask — examine and change blocked signals

35447 **SYNOPSIS**

35448 #include &lt;signal.h&gt;

35449 THR int pthread\_sigmask(int how, const sigset\_t \*restrict set,  
35450 sigset\_t \*restrict oset);35451 int sigprocmask(int how, const sigset\_t \*restrict set,  
35452 sigset\_t \*restrict oset);35453 **DESCRIPTION**35454 THR The *pthread\_sigmask()* function is used to examine or change (or both) the calling thread's signal mask, regardless of the number of threads in the process. The effect shall be the same as described for *sigprocmask()*, without the restriction that the call be made in a single-threaded process.35458 In a single-threaded process, the *sigprocmask()* function allows the calling process to examine or change (or both) the signal mask of the calling thread.35460 If the argument *set* is not a null pointer, it points to a set of signals to be used to change the currently blocked set.35462 The argument *how* indicates the way in which the set is changed, and the application shall ensure it consists of one of the following values:35464 SIG\_BLOCK The resulting set shall be the union of the current set and the signal set pointed to by *set*.35466 SIG\_SETMASK The resulting set shall be the signal set pointed to by *set*.35467 SIG\_UNBLOCK The resulting set shall be the intersection of the current set and the complement of the signal set pointed to by *set*.35469 If the argument *oset* is not a null pointer, the previous mask is stored in the location pointed to by *oset*. If *set* is a null pointer, the value of the argument *how* is not significant and the process' signal mask is unchanged; thus the call can be used to enquire about currently blocked signals.35472 If there are any pending unblocked signals after the call to *sigprocmask()*, at least one of those signals shall be delivered before the call to *sigprocmask()* returns.

35474 It is not possible to block those signals which cannot be ignored. This shall be enforced by the system without causing an error to be indicated.

35476 If any of the SIGFPE, SIGILL, SIGSEGV, or SIGBUS signals are generated while they are blocked, the result is undefined, unless the signal was generated by the *kill()* function, the *sigqueue()* function, or the *raise()* function.35479 If *sigprocmask()* fails, the thread's signal mask is not changed.35480 The use of the *sigprocmask()* function is unspecified in a multi-threaded process.35481 **RETURN VALUE**35482 THR Upon successful completion *pthread\_sigmask()* shall return 0; otherwise, it shall return the corresponding error number.35484 Upon successful completion, *sigprocmask()* shall return 0; otherwise, -1 shall be returned, *errno* shall be set to indicate the error, and the process' signal mask shall be unchanged.

35486 **ERRORS**

35487 THR The *pthread\_sigmask()* and *sigprocmask()* functions shall fail if:

35488 [EINVAL] The value of the *how* argument is not equal to one of the defined values.

35489 THR The *pthread\_sigmask()* function shall not return an error code of [EINTR].

35490 **EXAMPLES**

35491 None.

35492 **APPLICATION USAGE**

35493 None.

35494 **RATIONALE**

35495 When a process' signal mask is changed in a signal-catching function that is installed by  
35496 *sigaction()*, the restoration of the signal mask on return from the signal-catching function  
35497 overrides that change (see *sigaction()*). If the signal-catching function was installed with  
35498 *signal()*, it is unspecified whether this occurs.

35499 See *kill()* for a discussion of the requirement on delivery of signals.

35500 **FUTURE DIRECTIONS**

35501 None.

35502 **SEE ALSO**

35503 *sigaction()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*, *sigpending()*,  
35504 *sigqueue()*, *sigsuspend()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <signal.h>

35505 **CHANGE HISTORY**

35506 First released in Issue 3.

35507 Entry included for alignment with the POSIX.1-1988 standard.

35508 **Issue 4**

35509 The DESCRIPTION is changed to indicate that signals can also be generated by *raise()*.

35510 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 35511 • The type of the arguments *set* and *oset* are changed from **sigset\_t\*** to **const sigset\_t\***.

35512 **Issue 5**

35513 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

35514 The *pthread\_sigmask()* function is added for alignment with the POSIX Threads Extension.

35515 **Issue 6**

35516 The *pthread\_sigmask()* function is marked as part of the Threads option.

35517 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 35518 • The DESCRIPTION is updated to explicitly state the functions which may generate the  
35519 signal.

35520 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

35521 The **restrict** keyword is added to the *pthread\_sigmask()* and *sigprocmask()* prototypes for  
35522 alignment with the ISO/IEC 9899: 1999 standard.

35523 **NAME**

35524 pthread\_spin\_destroy, pthread\_spin\_init — destroy or initialize a spin lock object

35525 **SYNOPSIS**

35526 SPI #include &lt;pthread.h&gt;

35527 int pthread\_spin\_destroy(pthread\_spinlock\_t \*lock);

35528 int pthread\_spin\_init(pthread\_spinlock\_t \*lock, int pshared);

35529

35530 **DESCRIPTION**

35531 The *pthread\_spin\_destroy()* function destroys the spin lock referenced by *lock* and releases any  
 35532 resources used by the lock. The effect of subsequent use of the lock is undefined until the lock is  
 35533 reinitialized by another call to *pthread\_spin\_init()*. The results are undefined if  
 35534 *pthread\_spin\_destroy()* is called when a thread holds the lock, or if this function is called with an  
 35535 uninitialized thread spin lock.

35536 The *pthread\_spin\_init()* function allocates any resources required to use the spin lock referenced  
 35537 by *lock* and initializes the lock to an unlocked state.

35538 TSH If the Thread Process-Shared Synchronization option is supported and the value of *pshared* is  
 35539 PTHREAD\_PROCESS\_SHARED, the implementation shall permit the spin lock to be operated  
 35540 upon by any thread that has access to the memory where the spin lock is allocated, even if it is  
 35541 allocated in memory that is shared by multiple processes.

35542 If the Thread Process-Shared Synchronization option is supported and the value of *pshared* is  
 35543 PTHREAD\_PROCESS\_PRIVATE, or if the option is not supported, the spin lock shall only be  
 35544 operated upon by threads created within the same process as the thread that initialized the spin  
 35545 lock. If threads of differing processes attempt to operate on such a spin lock, the behavior is  
 35546 undefined.

35547 The results are undefined if *pthread\_spin\_init()* is called specifying an already initialized spin  
 35548 lock. The results are undefined if a spin lock is used without first being initialized.

35549 If the *pthread\_spin\_init()* function fails, the lock is not initialized and the contents of *lock* are  
 35550 undefined.

35551 Only the object referenced by *lock* may be used for performing synchronization.

35552 The result of referring to copies of that object in calls to *pthread\_spin\_destroy()*,  
 35553 *pthread\_spin\_lock()*, *pthread\_spin\_trylock()*, or *pthread\_spin\_unlock()* is undefined.

35554 **RETURN VALUE**

35555 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
 35556 be returned to indicate the error.

35557 **ERRORS**

35558 These functions may fail if:

35559 [EBUSY] The implementation has detected an attempt to initialize or destroy a spin  
 35560 lock while it is in use (for example, while being used in a *pthread\_spin\_lock()*  
 35561 call) by another thread.

35562 [EINVAL] The value specified by *lock* is invalid.

35563 The *pthread\_spin\_init()* function shall fail if:

35564 [EAGAIN] The system lacks the necessary resources to initialize another spin lock.

35565 [ENOMEM] Insufficient memory exists to initialize the lock.

35566 These functions shall not return an error code of [EINTR].

35567 **EXAMPLES**

35568 None.

35569 **APPLICATION USAGE**

35570 The *pthread\_spin\_destroy()* and *pthread\_spin\_init()* functions are part of the Spin Locks option  
35571 and need not be provided on all implementations.

35572 **RATIONALE**

35573 None.

35574 **FUTURE DIRECTIONS**

35575 None.

35576 **SEE ALSO**

35577 *pthread\_spin\_lock()*, *pthread\_spin\_trylock()*, *pthread\_spin\_unlock()*, the Base Definitions volume of  
35578 IEEE Std. 1003.1-200x, <<pthread.h>>

35579 **CHANGE HISTORY**

35580 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

35581 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

35582 **NAME**

35583 pthread\_spin\_init — initialize a spin lock object

35584 **SYNOPSIS**

35585 SPI #include &lt;pthread.h&gt;

35586 int pthread\_spin\_init(pthread\_spinlock\_t \*lock, int pshared);

35587

35588 **DESCRIPTION**35589 Refer to *pthread\_spin\_destroy()*.

35590 **NAME**

35591 pthread\_spin\_lock, pthread\_spin\_trylock — lock a spin lock object

35592 **SYNOPSIS**

35593 SPI #include &lt;pthread.h&gt;

35594 int pthread\_spin\_lock(pthread\_spinlock\_t \*lock);

35595 int pthread\_spin\_trylock(pthread\_spinlock\_t \*lock);

35596

35597 **DESCRIPTION**

35598 The *pthread\_spin\_lock()* function locks the spin lock referenced by *lock*. The calling thread  
35599 acquires the lock if it is not held by another thread. Otherwise, the thread spins (that is, does not  
35600 return from the *pthread\_spin\_lock()* call) until the lock becomes available. The results are  
35601 undefined if the calling thread holds the lock at the time the call is made. The  
35602 *pthread\_spin\_trylock()* function locks the spin lock referenced by *lock* if it is not held by any  
35603 thread. Otherwise, the function fails.

35604 The results are undefined if any of these functions is called with an uninitialized spin lock.

35605 **RETURN VALUE**

35606 Upon successful completion, these functions shall return zero; otherwise, an error number shall  
35607 be returned to indicate the error.

35608 **ERRORS**

35609 These functions may fail if:

35610 [EINVAL] The value specified by *lock* does not refer to an initialized spin lock object.35611 The *pthread\_spin\_lock()* function may fail if:

35612 [EDEADLK] The calling thread already holds the lock.

35613 The *pthread\_spin\_trylock()* function shall fail if:

35614 [EBUSY] A thread currently holds the lock.

35615 These functions shall not return an error code of [EINTR].

35616 **EXAMPLES**

35617 None.

35618 **APPLICATION USAGE**

35619 Applications using this function may be subject to priority inversion, as discussed in the Base  
35620 Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

35621 The *pthread\_spin\_lock()* and *pthread\_spin\_trylock()* functions are part of the Spin Locks option  
35622 and need not be provided on all implementations.

35623 **RATIONALE**

35624 None.

35625 **FUTURE DIRECTIONS**

35626 None.

35627 **SEE ALSO**

35628 *pthread\_spin\_init()*, *pthread\_spin\_destroy()*, *pthread\_spin\_unlock()*, the Base Definitions volume of  
35629 IEEE Std. 1003.1-200x, <pthread.h>

35630 **CHANGE HISTORY**

35631 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

35632 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

35633 **NAME**

35634 pthread\_spin\_trylock — lock a spin lock object

35635 **SYNOPSIS**

35636 SPI #include <pthread.h>

35637 int pthread\_spin\_trylock(pthread\_spinlock\_t \*lock);

35638

35639 **DESCRIPTION**

35640 Refer to *pthread\_spin\_lock()*.

35641 **NAME**

35642 pthread\_spin\_unlock — unlock a spin lock object

35643 **SYNOPSIS**

35644 SPI #include &lt;pthread.h&gt;

35645 int pthread\_spin\_unlock(pthread\_spinlock\_t \*lock);

35646

35647 **DESCRIPTION**

35648 The *pthread\_spin\_unlock()* function releases the spin lock referenced by *lock* which was locked  
35649 via the *pthread\_spin\_lock()* or *pthread\_spin\_trylock()* functions. The results are undefined if the  
35650 lock is not held by the calling thread. If there are threads spinning on the lock when  
35651 *pthread\_spin\_unlock()* is called, the lock becomes available and an unspecified spinning thread  
35652 shall acquire the lock.

35653 The results are undefined if this function is called with an uninitialized thread spin lock.

35654 **RETURN VALUE**

35655 Upon successful completion, the *pthread\_spin\_unlock()* function shall return zero; otherwise, an  
35656 error number shall be returned to indicate the error.

35657 **ERRORS**35658 The *pthread\_spin\_unlock()* function may fail if:

35659 [EINVAL] An invalid argument was specified.

35660 [EPERM] The calling thread does not hold the lock.

35661 This function shall not return an error code of [EINTR].

35662 **EXAMPLES**

35663 None.

35664 **APPLICATION USAGE**

35665 The *pthread\_spin\_unlock()* function is part of the Spin Locks option and need not be provided on  
35666 all implementations.

35667 **RATIONALE**

35668 None.

35669 **FUTURE DIRECTIONS**

35670 None.

35671 **SEE ALSO**

35672 *pthread\_spin\_init()*, *pthread\_spin\_destroy()*, *pthread\_spin\_lock()*, *pthread\_spin\_trylock()*, the Base  
35673 Definitions volume of IEEE Std. 1003.1-200x, <pthread.h>

35674 **CHANGE HISTORY**

35675 First released in Issue 6. Derived from IEEE Std. 1003.1j-2000.

35676 In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

35677 **NAME**

35678 pthread\_testcancel — set cancelability state

35679 **SYNOPSIS**

35680 THR #include <pthread.h>

35681 void pthread\_testcancel(void);

35682

35683 **DESCRIPTION**

35684 Refer to *pthread\_setcancelstate()*.

35685 **NAME**

35686 ptsname — get name of the slave pseudo-terminal device

35687 **SYNOPSIS**

35688 XSI #include &lt;stdlib.h&gt;

35689 char \*ptsname(int *fildev*);

35690

35691 **DESCRIPTION**

35692 The *ptsname()* function shall return the name of the slave pseudo-terminal device associated  
35693 with a master pseudo-terminal device. The *fildev* argument is a file descriptor that refers to the  
35694 master device. The *ptsname()* function shall return a pointer to a string containing the path name  
35695 of the corresponding slave device.

35696 The *ptsname()* function need not be reentrant. A function that is not required to be reentrant is  
35697 not required to be thread-safe.

35698 **RETURN VALUE**

35699 Upon successful completion, *ptsname()* shall return a pointer to a string which is the name of the  
35700 pseudo-terminal slave device. Upon failure, *ptsname()* shall return a null pointer. This could  
35701 occur if *fildev* is an invalid file descriptor or if the slave device name does not exist in the file  
35702 system.

35703 **ERRORS**

35704 No errors are defined.

35705 **EXAMPLES**

35706 None.

35707 **APPLICATION USAGE**35708 The value returned may point to a static data area that is overwritten by each call to *ptsname()*.35709 **RATIONALE**

35710 None.

35711 **FUTURE DIRECTIONS**

35712 None.

35713 **SEE ALSO**

35714 *grantpt()*, *open()*, *ttyname()*, *unlockpt()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
35715 <stdlib.h>

35716 **CHANGE HISTORY**

35717 First released in Issue 4, Version 2.

35718 **Issue 5**

35719 Moved from X/OPEN UNIX extension to BASE.

35720 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

35721 **NAME**

35722       putc — put byte on a stream

35723 **SYNOPSIS**

35724       #include &lt;stdio.h&gt;

35725       int putc(int *c*, FILE \**stream*);35726 **DESCRIPTION**

35727 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
35728       conflict between the requirements described here and the ISO C standard is unintentional. This  
35729       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

35730       The *putc()* function shall be equivalent to *fputc()*, except that if it is implemented as a macro it  
35731       may evaluate *stream* more than once, so the argument should never be an expression with side  
35732       effects.

35733 **RETURN VALUE**35734       Refer to *fputc()*.35735 **ERRORS**35736       Refer to *fputc()*.35737 **EXAMPLES**

35738       None.

35739 **APPLICATION USAGE**

35740       Because it may be implemented as a macro, *putc()* may treat a *stream* argument with side effects  
35741       incorrectly. In particular, *putc(c,\*f++)* does not necessarily work correctly. Therefore, use of this  
35742       function is not recommended in such situations; *fputc()* should be used instead.

35743 **RATIONALE**

35744       None.

35745 **FUTURE DIRECTIONS**

35746       None.

35747 **SEE ALSO**35748       *fputc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>35749 **CHANGE HISTORY**

35750       First released in Issue 1. Derived from Issue 1 of the SVID.

35751 **Issue 4**

35752       The APPLICATION USAGE section now states that the use of this function is not recommended  
35753       with a *stream* argument with side effects.

35754       The following change is incorporated for alignment with the ISO C standard:

- 35755       • The *c* argument is not allowed to be evaluated more than once.

35756 **NAME**35757        `putc_unlocked` — stdio with explicit client locking35758 **SYNOPSIS**35759 TSF        `#include <stdio.h>`35760        `int putc_unlocked(int c, FILE *stream);`

35761

35762 **DESCRIPTION**35763        Refer to `getc_unlocked()`.

35764 **NAME**

35765 putchar — put byte on stdout stream

35766 **SYNOPSIS**

35767 #include <stdio.h>

35768 int putchar(int c);

35769 **DESCRIPTION**

35770 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
35771 conflict between the requirements described here and the ISO C standard is unintentional. This  
35772 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

35773 The function call *putchar(c)* shall be equivalent to *putc(c,stdout)*.

35774 **RETURN VALUE**

35775 Refer to *fputc()*.

35776 **ERRORS**

35777 Refer to *fputc()*.

35778 **EXAMPLES**

35779 None.

35780 **APPLICATION USAGE**

35781 None.

35782 **RATIONALE**

35783 None.

35784 **FUTURE DIRECTIONS**

35785 None.

35786 **SEE ALSO**

35787 *putc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>

35788 **CHANGE HISTORY**

35789 First released in Issue 1. Derived from Issue 1 of the SVID.

35790 **NAME**

35791 putchar\_unlocked — stdio with explicit client locking

35792 **SYNOPSIS**

35793 TSF #include &lt;stdio.h&gt;

35794 int putchar\_unlocked(int c);

35795

35796 **DESCRIPTION**35797 Refer to *getc\_unlocked()*.

35798 **NAME**

35799 putenv — change or add a value to environment

35800 **SYNOPSIS**

35801 xSI #include &lt;stdlib.h&gt;

35802 int putenv(char \*string);

35803

35804 **DESCRIPTION**

35805 The *putenv()* function uses the *string* argument to set environment variable values. The *string*  
35806 argument should point to a string of the form "*name=value*". The *putenv()* function makes the  
35807 value of the environment variable *name* equal to *value* by altering an existing variable or creating  
35808 a new one. In either case, the string pointed to by *string* becomes part of the environment, so  
35809 altering the string shall change the environment. The space used by *string* is no longer used once  
35810 a new string-defining *name* is passed to *putenv()*.

35811 The *putenv()* function need not be reentrant. A function that is not required to be reentrant is not  
35812 required to be thread-safe.

35813 **RETURN VALUE**

35814 Upon successful completion, *putenv()* shall return 0; otherwise, it shall return a non-zero value  
35815 and set *errno* to indicate the error.

35816 **ERRORS**35817 The *putenv()* function may fail if:

35818 [ENOMEM] Insufficient memory was available.

35819 **EXAMPLES**35820 **Changing the Value of an Environment Variable**

35821 The following example changes the value of the *HOME* environment variable to the value  
35822 */usr/home*.

35823 #include &lt;stdlib.h&gt;

35824 ...

35825 static char \*var = "HOME=/usr/home";

35826 int ret;

35827 ret = putenv(var);

35828 **APPLICATION USAGE**

35829 The *putenv()* function manipulates the environment pointed to by *environ*, and can be used in  
35830 conjunction with *getenv()*.

35831 This routine may use *malloc()* to enlarge the environment.

35832 A potential error is to call *putenv()* with an automatic variable as the argument, then return from  
35833 the calling function while *string* is still part of the environment.

35834 The *setenv()* function is preferred over this function.35835 **RATIONALE**

35836 None.

35837 **FUTURE DIRECTIONS**

35838 None.

35839 **SEE ALSO**35840 *exec*, *getenv()*, *malloc()*, *setenv()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h> |35841 **CHANGE HISTORY**

35842 First released in Issue 1. Derived from Issue 1 of the SVID. |

35843 **Issue 4**

35844 The &lt;stdlib.h&gt; header is added to the SYNOPSIS section.

35845 The type of argument *string* is changed from **char\*** to **const char\***.35846 **Issue 5**35847 The type of the argument to this function is changed from **const char\*** to **char\***. This was indicated as a FUTURE DIRECTION in previous issues.

35849 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

## 35850 NAME

35851 putmsg, putpmsg — send a message on a STREAM (STREAMS)

## 35852 SYNOPSIS

35853 XSR #include &lt;stropts.h&gt;

```

35854 int putmsg(int fildev, const struct strbuf *ctlptr,
35855 const struct strbuf *dataptr, int flags);
35856 int putpmsg(int fildev, const struct strbuf *ctlptr,
35857 const struct strbuf *dataptr, int band, int flags);
35858

```

## 35859 DESCRIPTION

35860 The *putmsg()* function shall create a message from a process buffer(s) and send the message to a  
 35861 STREAMS file. The message may contain either a data part, a control part, or both. The data and  
 35862 control parts are distinguished by placement in separate buffers, as described below. The  
 35863 semantics of each part are defined by the STREAMS module that receives the message.

35864 The *putpmsg()* function does the same thing as *putmsg()*, but the process can send messages in  
 35865 different priority bands. Except where noted, all requirements on *putmsg()* also pertain to  
 35866 *putpmsg()*.

35867 The *fildev* argument specifies a file descriptor referencing an open STREAM. The *ctlptr* and  
 35868 *dataptr* arguments each point to a **strbuf** structure.

35869 The *ctlptr* argument points to the structure describing the control part, if any, to be included in  
 35870 the message. The *buf* member in the **strbuf** structure points to the buffer where the control  
 35871 information resides, and the *len* member indicates the number of bytes to be sent. The *maxlen*  
 35872 member is not used by *putmsg()*. In a similar manner, the argument *dataptr* specifies the data, if  
 35873 any, to be included in the message. The *flags* argument indicates what type of message should be  
 35874 sent and is described further below.

35875 To send the data part of a message, the application shall ensure that *dataptr* is not a null pointer  
 35876 and the *len* member of *dataptr* is 0 or greater. To send the control part of a message, the  
 35877 application shall ensure that the corresponding values are set for *ctlptr*. No data (control) part  
 35878 shall be sent if either *dataptr(ctlptr)* is a null pointer or the *len* member of *dataptr(ctlptr)* is set to  
 35879 -1.

35880 For *putmsg()*, if a control part is specified and *flags* is set to RS\_HIPRI, a high priority message is  
 35881 sent. If no control part is specified, and *flags* is set to RS\_HIPRI, *putmsg()* fails and sets *errno* to  
 35882 [EINVAL]. If *flags* is set to 0, a normal message (priority band equal to 0) is sent. If a control part  
 35883 and data part are not specified and *flags* is set to 0, no message is sent and 0 is returned.

35884 For *putpmsg()*, the flags are different. The *flags* argument is a bitmask with the following  
 35885 mutually-exclusive flags defined: MSG\_HIPRI and MSG\_BAND. If *flags* is set to 0, *putpmsg()*  
 35886 fails and sets *errno* to [EINVAL]. If a control part is specified and *flags* is set to MSG\_HIPRI and  
 35887 *band* is set to 0, a high-priority message is sent. If *flags* is set to MSG\_HIPRI and either no control  
 35888 part is specified or *band* is set to a non-zero value, *putpmsg()* fails and sets *errno* to [EINVAL]. If  
 35889 *flags* is set to MSG\_BAND, then a message is sent in the priority band specified by *band*. If a  
 35890 control part and data part are not specified and *flags* is set to MSG\_BAND, no message is sent  
 35891 and 0 is returned.

35892 The *putmsg()* function blocks if the STREAM write queue is full due to internal flow control  
 35893 conditions, with the following exceptions:

- 35894 • For high-priority messages, *putmsg()* does not block on this condition and continues  
 35895 processing the message.

35896           • For other messages, *putmsg()* does not block but fails when the write queue is full and  
35897           O\_NONBLOCK is set.

35898           The *putmsg()* function also blocks, unless prevented by lack of internal resources, while waiting  
35899           for the availability of message blocks in the STREAM, regardless of priority or whether  
35900           O\_NONBLOCK has been specified. No partial message is sent.

#### 35901 RETURN VALUE

35902           Upon successful completion, *putmsg()* and *putpmsg()* shall return 0; otherwise, they shall return  
35903           -1 and set *errno* to indicate the error.

#### 35904 ERRORS

35905           The *putmsg()* and *putpmsg()* functions shall fail if:

35906           [EAGAIN]           A non-priority message was specified, the O\_NONBLOCK flag is set, and the  
35907           STREAM write queue is full due to internal flow control conditions; or buffers  
35908           could not be allocated for the message that was to be created.

35909           [EBADF]           *fildes* is not a valid file descriptor open for writing.

35910           [EINTR]           A signal was caught during *putmsg()*.

35911           [EINVAL]           An undefined value is specified in *flags*, or *flags* is set to RS\_HIPRI or  
35912           MSG\_HIPRI and no control part is supplied, or the STREAM or multiplexer  
35913           referenced by *fildes* is linked (directly or indirectly) downstream from a  
35914           multiplexer, or *flags* is set to MSG\_HIPRI and *band* is non-zero (for *putpmsg()*  
35915           only).

35916           [ENOSR]           Buffers could not be allocated for the message that was to be created due to  
35917           insufficient STREAMS memory resources.

35918           [ENOSTR]          A STREAM is not associated with *fildes*.

35919           [ENXIO]           A hangup condition was generated downstream for the specified STREAM.

35920           [EPIPE] or [EIO]   The *fildes* argument refers to a STREAMS-based pipe and the other end of the  
35921           pipe is closed. A SIGPIPE signal is generated for the calling thread.

35922           [ERANGE]          The size of the data part of the message does not fall within the range  
35923           specified by the maximum and minimum packet sizes of the topmost  
35924           STREAM module. This value is also returned if the control part of the message  
35925           is larger than the maximum configured size of the control part of a message,  
35926           or if the data part of a message is larger than the maximum configured size of  
35927           the data part of a message.

35928           In addition, *putmsg()* and *putpmsg()* shall fail if the STREAM head had processed an  
35929           asynchronous error before the call. In this case, the value of *errno* does not reflect the result of  
35930           *putmsg()* or *putpmsg()*, but reflects the prior error.

35931 **EXAMPLES**35932 **Sending a High-Priority Message**

35933 The value of *fd* is assumed to refer to an open STREAMS file. This call to *putmsg()* does the  
35934 following:

- 35935 1. Creates a high-priority message with a control part and a data part, using the buffers  
35936 pointed to by *ctrlbuf* and *databuf*, respectively.
- 35937 2. Sends the message to the STREAMS file identified by *fd*.

```
35938 #include <stropts.h>
35939 #include <string.h>
35940 ...
35941 int fd;
35942 char *ctrlbuf = "This is the control part";
35943 char *databuf = "This is the data part";
35944 struct strbuf ctrl;
35945 struct strbuf data;
35946 int ret;

35947 ctrl.buf = ctrlbuf;
35948 ctrl.len = strlen(ctrlbuf);

35949 data.buf = databuf;
35950 data.len = strlen(databuf);

35951 ret = putmsg(fd, &ctrl, &data, MSG_HIPRI);
```

35952 **Using putpmsg()**

35953 This example has the same effect as the previous example. In this example, however, the  
35954 *putpmsg()* function is used to create and send the message to the STREAMS file.

```
35955 #include <stropts.h>
35956 #include <string.h>
35957 ...
35958 int fd;
35959 char *ctrlbuf = "This is the control part";
35960 char *databuf = "This is the data part";
35961 struct strbuf ctrl;
35962 struct strbuf data;
35963 int ret;

35964 ctrl.buf = ctrlbuf;
35965 ctrl.len = strlen(ctrlbuf);

35966 data.buf = databuf;
35967 data.len = strlen(databuf);

35968 ret = putpmsg(fd, &ctrl, &data, 0, MSG_HIPRI);
```

35969 **APPLICATION USAGE**

35970 None.

35971 **RATIONALE**

35972 None.

35973 **FUTURE DIRECTIONS**

35974 None.

35975 **SEE ALSO**35976 *getmsg()*, *poll()*, *read()*, *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
35977 **<stropts.h>**, Section 2.6 (on page 539)35978 **CHANGE HISTORY**

35979 First released in Issue 4, Version 2.

35980 **Issue 5**

35981 Moved from X/OPEN UNIX extension to BASE.

35982 The following text is removed from the DESCRIPTION: “The STREAM head guarantees that the  
35983 control part of a message generated by *putmsg()* is at least 64 bytes in length”.35984 **Issue 6**

35985 This function is marked as part of the XSI STREAMS Option Group.

35986 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

35987 **NAME**

35988 puts — put a string on standard output

35989 **SYNOPSIS**

35990 #include &lt;stdio.h&gt;

35991 int puts(const char \*s);

35992 **DESCRIPTION**

35993 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 35994 conflict between the requirements described here and the ISO C standard is unintentional. This  
 35995 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

35996 The *puts()* function shall write the string pointed to by *s*, followed by a <newline> character, to  
 35997 the standard output stream *stdout*. The terminating null byte shall not be written.

35998 cx The *st\_ctime* and *st\_mtime* fields of the file shall be marked for update between the successful  
 35999 execution of *puts()* and the next successful completion of a call to *fflush()* or *fclose()* on the same  
 36000 stream or a call to *exit()* or *abort()*.

36001 **RETURN VALUE**

36002 Upon successful completion, *puts()* shall return a non-negative number. Otherwise, it shall  
 36003 cx return EOF, shall set an error indicator for the stream, and *errno* shall be set to indicate the error.

36004 **ERRORS**36005 Refer to *fputc()*.36006 **EXAMPLES**36007 **Printing to Standard Output**

36008 The following example gets the current time, converts it to a string using *localtime()* and  
 36009 *asctime()*, and prints it to standard output using *puts()*. It then prints the number of minutes to  
 36010 an event for which it is waiting.

```
36011 #include <time.h>
36012 #include <stdio.h>
36013 ...
36014 time_t now;
36015 int minutes_to_event;
36016 ...
36017 time(&now);
36018 printf("The time is ");
36019 puts(asctime(localtime(&now)));
36020 printf("There are %d minutes to the event.\n",
36021 minutes_to_event);
36022 ...
```

36023 **APPLICATION USAGE**36024 The *puts()* function appends a <newline> character, while *fputs()* does not.36025 **RATIONALE**

36026 None.

36027 **FUTURE DIRECTIONS**

36028 None.

36029 **SEE ALSO**

36030 *fopen()*, *fputs()*, *putc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>

36031 **CHANGE HISTORY**

36032 First released in Issue 1. Derived from Issue 1 of the SVID.

36033 **Issue 4**

36034 In the DESCRIPTION, the words “null character” are replaced by “null byte”.

36035 The following change is incorporated for alignment with the ISO C standard:

- 36036 • The type of argument *s* is changed from **char\*** to **const char\***.

36037 **Issue 6**

36038 Extensions beyond the ISO C standard are now marked.

36039 **NAME**

36040 pututxline — put an entry into user accounting database

36041 **SYNOPSIS**

36042 xSI `#include <utmpx.h>`

36043 `struct utmpx *pututxline(const struct utmpx *utmpx);`

36044

36045 **DESCRIPTION**

36046 Refer to *endutxent()*.

36047 **NAME**

36048 putwc — put a wide character on a stream

36049 **SYNOPSIS**

36050 #include &lt;stdio.h&gt;

36051 #include &lt;wchar.h&gt;

36052 wint\_t putwc(wchar\_t *wc*, FILE \**stream*);36053 **DESCRIPTION**

36054 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
36055 conflict between the requirements described here and the ISO C standard is unintentional. This  
36056 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

36057 The *putwc()* function shall be equivalent to *fputwc()*, except that if it is implemented as a macro  
36058 it may evaluate *stream* more than once, so the argument should never be an expression with side  
36059 effects.

36060 **RETURN VALUE**36061 Refer to *fputwc()*.36062 **ERRORS**36063 Refer to *fputwc()*.36064 **EXAMPLES**

36065 None.

36066 **APPLICATION USAGE**

36067 Because it may be implemented as a macro, *putwc()* may treat a *stream* argument with side  
36068 effects incorrectly. In particular, *putwc(wc,\*f++)* need not work correctly. Therefore, use of this  
36069 function is not recommended; *fputwc()* should be used instead.

36070 **RATIONALE**

36071 None.

36072 **FUTURE DIRECTIONS**

36073 None.

36074 **SEE ALSO**36075 *fputwc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>, <wchar.h>36076 **CHANGE HISTORY**

36077 First released as a World-wide Portability Interface in Issue 4.

36078 **Issue 5**

36079 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the type of argument *wc*  
36080 is changed from **wint\_t** to **wchar\_t**.

36081 The Optional Header (OH) marking is removed from &lt;stdio.h&gt;.

36082 **NAME**

36083 putwchar — put a wide character on stdout stream

36084 **SYNOPSIS**

36085 #include <wchar.h>

36086 wint\_t putwchar(wchar\_t wc);

36087 **DESCRIPTION**

36088 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
36089 conflict between the requirements described here and the ISO C standard is unintentional. This  
36090 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

36091 The function call *putwchar(wc)* shall be equivalent to *putwc(wc,stdout)*.

36092 **RETURN VALUE**

36093 Refer to *fputwc()*.

36094 **ERRORS**

36095 Refer to *fputwc()*.

36096 **EXAMPLES**

36097 None.

36098 **APPLICATION USAGE**

36099 None.

36100 **RATIONALE**

36101 None.

36102 **FUTURE DIRECTIONS**

36103 None.

36104 **SEE ALSO**

36105 *fputwc()*, *putwc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <wchar.h>

36106 **CHANGE HISTORY**

36107 First released in Issue 4.

36108 **Issue 5**

36109 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the type of argument *wc*  
36110 is changed from **wint\_t** to **wchar\_t**.

36111 **NAME**

36112            pwrite, — write on a file

36113 **SYNOPSIS**

36114            #include &lt;unistd.h&gt;

36115 XSI        ssize\_t pwrite(int *fildev*, const void \**buf*, size\_t *nbyte*,  
36116               off\_t *offset*);

36117

36118 **DESCRIPTION**36119            Refer to *write()*.

36120 **NAME**

36121 qsort — sort a table of data

36122 **SYNOPSIS**

36123 #include &lt;stdlib.h&gt;

36124 void qsort(void \*base, size\_t nel, size\_t width  
36125 int (\*compar)(const void \*, const void \*));36126 **DESCRIPTION**36127 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
36128 conflict between the requirements described here and the ISO C standard is unintentional. This  
36129 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.36130 The *qsort()* function sorts an array of *nel* objects, the initial element of which is pointed to by  
36131 *base*. The size of each object, in bytes, is specified by the *width* argument.36132 The contents of the array are sorted in ascending order according to a comparison function. The  
36133 *compar* argument is a pointer to the comparison function, which is called with two arguments  
36134 that point to the elements being compared. The application shall ensure that the function returns  
36135 an integer less than, equal to, or greater than 0, if the first argument is considered respectively  
36136 less than, equal to, or greater than the second. If two members compare as equal, their order in  
36137 the sorted array is unspecified.36138 **RETURN VALUE**36139 The *qsort()* function shall return no value.36140 **ERRORS**

36141 No errors are defined.

36142 **EXAMPLES**

36143 None.

36144 **APPLICATION USAGE**36145 The comparison function need not compare every byte, so arbitrary data may be contained in  
36146 the elements in addition to the values being compared.36147 **RATIONALE**

36148 None.

36149 **FUTURE DIRECTIONS**

36150 None.

36151 **SEE ALSO**

36152 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;stdlib.h&gt;

36153 **CHANGE HISTORY**

36154 First released in Issue 1. Derived from Issue 1 of the SVID.

36155 **Issue 4**

36156 The following change is incorporated for alignment with the ISO C standard:

- 36157
- The arguments to *compar* are formally defined in the SYNOPSIS section.

36158 **Issue 6**

36159 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

36160 **NAME**

36161 raise — send a signal to the executing process

36162 **SYNOPSIS**

36163 #include &lt;signal.h&gt;

36164 int raise(int *sig*);36165 **DESCRIPTION**

36166 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 36167 conflict between the requirements described here and the ISO C standard is unintentional. This  
 36168 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

36169 CX The *raise()* function shall send the signal *sig* to the executing thread or process. If a signal  
 36170 handler is called, the *raise()* function shall not return until after the signal handler does.

36171 THR If the implementation supports the Threads option, the effect of the *raise()* function is equivalent  
 36172 to calling:

36173 pthread\_kill(pthread\_self(), sig);

36174

36175 CX Otherwise, the effect of the *raise()* function is equivalent to calling:

36176 kill(getpid(), sig);

36177

36178 **RETURN VALUE**

36179 CX Upon successful completion, 0 shall be returned. Otherwise, a non-zero value shall be returned  
 36180 and *errno* shall be set to indicate the error.

36181 **ERRORS**36182 The *raise()* function shall fail if:36183 CX [EINVAL] The value of the *sig* argument is an invalid signal number.36184 **EXAMPLES**

36185 None.

36186 **APPLICATION USAGE**

36187 None.

36188 **RATIONALE**

36189 The term “thread” is an extension to the ISO C standard.

36190 **FUTURE DIRECTIONS**

36191 None.

36192 **SEE ALSO**

36193 *kill()*, *sigaction()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <signal.h>,  
 36194 <sys/types.h>

36195 **CHANGE HISTORY**

36196 First released in Issue 4. Derived from the ANSI C standard.

36197 **Issue 5**

36198 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

36199 **Issue 6**

36200 Extensions beyond the ISO C standard are now marked.

36201 The following new requirements on POSIX implementations derive from alignment with the  
36202 Single UNIX Specification:

- 36203 • In the RETURN VALUE section, the requirement to set *errno* on error is added.
- 36204 • The [EINVAL] error condition is added.

36205 **NAME**

36206 rand, rand\_r, srand — pseudo-random number generator

36207 **SYNOPSIS**

```
36208 #include <stdlib.h>
36209 int rand(void);
36210 TSF int rand_r(unsigned *seed);
36211 void srand(unsigned seed);
```

36212 **DESCRIPTION**

36213 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 36214 conflict between the requirements described here and the ISO C standard is unintentional. This  
 36215 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

36216 The *rand()* function shall compute a sequence of pseudo-random integers in the range 0 to  
 36217 XSI {RAND\_MAX} with a period of at least  $2^{32}$ .

36218 CX The *rand()* function need not be reentrant. A function that is not required to be reentrant is not  
 36219 required to be thread-safe.

36220 TSF The *rand\_r()* function shall compute a sequence of pseudo-random integers in the range 0 to  
 36221 {RAND\_MAX}. (The value of the {RAND\_MAX} macro shall be at least 32 767.)

36222 If *rand\_r()* is called with the same initial value for the object pointed to by *seed* and that object is  
 36223 not modified between successive returns and calls to *rand\_r()*, the same sequence shall be  
 36224 generated.

36225 The *srand()* function uses the argument as a seed for a new sequence of pseudo-random  
 36226 numbers to be returned by subsequent calls to *rand()*. If *srand()* is then called with the same  
 36227 seed value, the sequence of pseudo-random numbers shall be repeated. If *rand()* is called before  
 36228 any calls to *srand()* are made, the same sequence shall be generated as when *srand()* is first  
 36229 called with a seed value of 1.

36230 The implementation shall behave as if no function defined in this volume of  
 36231 IEEE Std. 1003.1-200x calls *rand()* or *srand()*.

36232 **RETURN VALUE**

36233 The *rand()* function shall return the next pseudo-random number in the sequence.

36234 TSF The *rand\_r()* function shall return a pseudo-random integer.

36235 The *srand()* function shall return no value.

36236 **ERRORS**

36237 No errors are defined.

36238 **EXAMPLES**36239 **Generating a Pseudo-Random Number Sequence**

36240 The following example demonstrates how to generate a sequence of pseudo-random numbers.

```
36241 #include <stdio.h>
36242 #include <stdlib.h>
36243 ...
36244 long count, i;
36245 char *keystri;
36246 int elementlen, len;
36247 char c;
```

```

36248 ...
36249 /* Initial random number generator. */
36250 srand(1);

36251 /* Create keys using only lower case characters */
36252 len = 0;
36253 for (i=0; i<count; i++) {
36254 while (len < elementlen) {
36255 c = (char) (rand() % 128);
36256 if (islower(c))
36257 keystr[len++] = c;
36258 }

36259 keystr[len] = '\0';
36260 printf("%s Element%0*ld\n", keystr, elementlen, i);
36261 len = 0;
36262 }

```

### 36263 **Generating the Same Sequence on Different Machines**

36264 The following code defines a pair of functions that could be incorporated into applications  
 36265 wishing to ensure that the same sequence of numbers is generated across different machines.

```

36266 static unsigned long next = 1;
36267 int myrand(void) /* RAND_MAX assumed to be 32767. */
36268 {
36269 next = next * 1103515245 + 12345;
36270 return((unsigned)(next/65536) % 32768);
36271 }

36272 void mysrand(unsigned seed)
36273 {
36274 next = seed;
36275 }

```

### 36276 **APPLICATION USAGE**

36277 The *drand48()* function provides a much more elaborate random number generator.

### 36278 **RATIONALE**

36279 The ISO C standard *rand()* and *srand()* functions allow per-process pseudo-random streams  
 36280 shared by all threads. Those two functions need not change, but there has to be mutual-  
 36281 exclusion that prevents interference between two threads concurrently accessing the random  
 36282 number generator.

36283 With regard to *rand()*, there are two different behaviors that may be wanted in a multi-threaded  
 36284 program:

- 36285 1. A single per-process sequence of pseudo-random numbers that is shared by all threads  
 36286 that call *rand()*
- 36287 2. A different sequence of pseudo-random numbers for each thread that calls *rand()*

36288 This is provided by the modified thread-safe function based on whether the seed value is global  
 36289 to the entire process or local to each thread.

36290 This does not address the known deficiencies of the *rand()* function implementations, which  
 36291 have been approached by maintaining more state. In effect, this specifies new thread-safe forms  
 36292 of a deficient function. Since alternatives to *rand()* are not standardized, they are not modified as

- 36293 part of this volume of IEEE Std. 1003.1-200x.
- 36294 **FUTURE DIRECTIONS**
- 36295 None.
- 36296 **SEE ALSO**
- 36297 *drand48()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>
- 36298 **CHANGE HISTORY**
- 36299 First released in Issue 1. Derived from Issue 1 of the SVID.
- 36300 **Issue 4**
- 36301 The definition of *srand()* is added to the SYNOPSIS section.
- 36302 In the DESCRIPTION, the text referring to the period of pseudo-random numbers is marked as  
36303 an extension.
- 36304 The example in the APPLICATION USAGE section is updated as follows:
- 36305 • To use ISO C standard syntax.
- 36306 • To avoid name clashes with standard functions.
- 36307 The following changes are incorporated for alignment with the ISO C standard:
- 36308 • The argument list of *rand()* is explicitly defined as **void**.
- 36309 • The argument *seed* is explicitly defined as **unsigned**.
- 36310 **Issue 5**
- 36311 The *rand\_r()* function is included for alignment with the POSIX Threads Extension.
- 36312 A note indicating that the *rand()* function need not be reentrant is added to the DESCRIPTION.
- 36313 **Issue 6**
- 36314 Extensions beyond the ISO C standard are now marked.
- 36315 The *rand\_r()* function is marked as part of the Thread-Safe Functions option.

36316 **NAME**

36317           random — generate pseudorandom number

36318 **SYNOPSIS**

36319 xSI       #include <stdlib.h>

36320           long random(void);

36321

36322 **DESCRIPTION**

36323           Refer to *initstate()*.

36324 **NAME**

36325 pread, read, readv — read from a file

36326 **SYNOPSIS**

36327 #include &lt;unistd.h&gt;

36328 XSI ssize\_t pread(int *fildes*, void \**buf*, size\_t *nbyte*, off\_t *offset*);36329 ssize\_t read(int *fildes*, void \**buf*, size\_t *nbyte*);

36330 XSI #include &lt;sys/uio.h&gt;

36331 ssize\_t readv(int *fildes*, const struct iovec \**iov*, int *iovcnt*);

36332

36333 **DESCRIPTION**

36334 The *read()* function attempts to read *nbyte* bytes from the file associated with the open file  
 36335 descriptor, *fildes*, into the buffer pointed to by *buf*. The behavior of multiple concurrent reads on  
 36336 the same pipe, FIFO, or terminal device is unspecified.

36337 If *nbyte* is zero, the *read()* function may detect and return errors as described below. In the  
 36338 absence of errors, or if error detection is not performed, the *read()* function shall return zero and  
 36339 have no other results.

36340 On files that support seeking (for example, a regular file), the *read()* starts at a position in the file  
 36341 given by the file offset associated with *fildes*. The file offset is incremented by the number of  
 36342 bytes actually read.

36343 Files that do not support seeking—for example, terminals—always read from the current  
 36344 position. The value of a file offset associated with such a file is undefined.

36345 No data transfer shall occur past the current end-of-file. If the starting position is at or after the  
 36346 end-of-file, 0 shall be returned. If the file refers to a device special file, the result of subsequent  
 36347 *read()* requests is implementation-defined.

36348 If the value of *nbyte* is greater than {SSIZE\_MAX}, the result is implementation-defined.

36349 When attempting to read from an empty pipe or FIFO:

- 36350 • If no process has the pipe open for writing, *read()* shall return 0 to indicate end-of-file.
- 36351 • If some process has the pipe open for writing and O\_NONBLOCK is set, *read()* shall return  
 36352 -1 and set *errno* to [EAGAIN].
- 36353 • If some process has the pipe open for writing and O\_NONBLOCK is clear, *read()* shall block  
 36354 the calling thread until some data is written or the pipe is closed by all processes that had the  
 36355 pipe open for writing.

36356 When attempting to read a file (other than a pipe or FIFO) that supports non-blocking reads and  
 36357 has no data currently available:

- 36358 • If O\_NONBLOCK is set, *read()* shall return -1 and set *errno* to [EAGAIN].
- 36359 • If O\_NONBLOCK is clear, *read()* shall block the calling thread until some data becomes  
 36360 available.
- 36361 • The use of the O\_NONBLOCK flag has no effect if there is some data available.

36362 The *read()* function reads data previously written to a file. If any portion of a regular file prior to  
 36363 the end-of-file has not been written, *read()* shall return bytes with value 0. For example, *lseek()*  
 36364 allows the file offset to be set beyond the end of existing data in the file. If data is later written at  
 36365 this point, subsequent reads in the gap between the previous end of data and the newly written  
 36366 data shall return bytes with value 0 until data is written into the gap.

36367 Upon successful completion, where *nbyte* is greater than 0, *read()* shall mark for update the  
36368 *st\_atime* field of the file, and shall return the number of bytes read. This number shall never be  
36369 greater than *nbyte*. The value returned may be less than *nbyte* if the number of bytes left in the  
36370 file is less than *nbyte*, if the *read()* request was interrupted by a signal, or if the file is a pipe or  
36371 FIFO or special file and has fewer than *nbyte* bytes immediately available for reading. For  
36372 example, a *read()* from a file associated with a terminal may return one typed line of data.

36373 If a *read()* is interrupted by a signal before it reads any data, it shall return  $-1$  with *errno* set to  
36374 [EINTR].

36375 If a *read()* is interrupted by a signal after it has successfully read some data, it shall return the  
36376 number of bytes read.

36377 XSR A *read()* from a STREAMS file can read data in three different modes: *byte-stream* mode,  
36378 *message-nondiscard* mode, and *message-discard* mode. The default is byte-stream mode. This can  
36379 be changed using the *I\_SRDOPT ioctl()* request, and can be tested with the *I\_GRDOPT ioctl()*. In  
36380 byte-stream mode, *read()* retrieves data from the STREAM until as many bytes as were  
36381 requested are transferred, or until there is no more data to be retrieved. Byte-stream mode  
36382 ignores message boundaries.

36383 In STREAMS message-nondiscard mode, *read()* retrieves data until as many bytes as were  
36384 requested are transferred, or until a message boundary is reached. If *read()* does not retrieve all  
36385 the data in a message, the remaining data is left on the STREAM, and can be retrieved by the  
36386 next *read()* call. Message-discard mode also retrieves data until as many bytes as were requested  
36387 are transferred, or a message boundary is reached. However, unread data remaining in a  
36388 message after the *read()* returns is discarded, and is not available for a subsequent *read()*,  
36389 *readv()*, or *getmsg()* call.

36390 How *read()* handles zero-byte STREAMS messages is determined by the current read mode  
36391 setting. In byte-stream mode, *read()* accepts data until it has read *nbyte* bytes, or until there is no  
36392 more data to read, or until a zero-byte message block is encountered. The *read()* function shall  
36393 then return the number of bytes read, and place the zero-byte message back on the STREAM to  
36394 be retrieved by the next *read()*, *readv()*, or *getmsg()*. In message-nondiscard mode or message-  
36395 discard mode, a zero-byte message shall return 0 and the message shall be removed from the  
36396 STREAM. When a zero-byte message is read as the first message on a STREAM, the message  
36397 shall be removed from the STREAM and 0 shall be returned, regardless of the read mode.

36398 A *read()* from a STREAMS file shall return the data in the message at the front of the STREAM  
36399 head read queue, regardless of the priority band of the message.

36400 By default, STREAMS are in control-normal mode, in which a *read()* from a STREAMS file can  
36401 only process messages that contain a data part but do not contain a control part. The *read()* shall  
36402 fail if a message containing a control part is encountered at the STREAM head. This default  
36403 action can be changed by placing the STREAM in either control-data mode or control-discard  
36404 mode with the *I\_SRDOPT ioctl()* command. In control-data mode, *read()* converts any control  
36405 part to data and passes it to the application before passing any data part originally present in the  
36406 same message. In control-discard mode, *read()* shall discard message control parts but return to  
36407 the process any data part in the message.

36408 In addition, *read()* and *readv()* fail if the STREAM head had processed an asynchronous error  
36409 before the call. In this case, the value of *errno* does not reflect the result of *read()* or *readv()*, but  
36410 reflects the prior error. If a hangup occurs on the STREAM being read, *read()* continues to  
36411 operate normally until the STREAM head read queue is empty. Thereafter, it shall return 0.

36412 XSI The *readv()* function is equivalent to *read()*, but places the input data into the *iovcnt* buffers  
36413 specified by the members of the *iov* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt*-1]. The *iovcnt* argument is  
36414 valid if greater than 0 and less than or equal to {IOV\_MAX}.

|       |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 36415 |     | Each <i>iovec</i> entry specifies the base address and length of an area in memory where data should be placed. The <i>readv()</i> function always fills an area completely before proceeding to the next.                                                                                                                                                                                                                             |
| 36416 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36417 |     | Upon successful completion, <i>readv()</i> shall mark for update the <i>st_atime</i> field of the file.                                                                                                                                                                                                                                                                                                                                |
| 36418 | SIO | If the <i>O_DSYNC</i> and <i>O_RSYNC</i> bits have been set, read I/O operations on the file descriptor complete as defined by synchronized I/O data integrity completion. If the <i>O_SYNC</i> and <i>O_RSYNC</i> bits have been set, read I/O operations on the file descriptor complete as defined by synchronized I/O file integrity completion.                                                                                   |
| 36419 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36420 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36421 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36422 | SHM | If <i>fildev</i> refers to a shared memory object, the result of the <i>read()</i> function is unspecified.                                                                                                                                                                                                                                                                                                                            |
| 36423 | TYM | If <i>fildev</i> refers to a typed memory object, the result of the <i>read()</i> function is unspecified.                                                                                                                                                                                                                                                                                                                             |
| 36424 |     | For regular files, no data transfer shall occur past the offset maximum established in the open file description associated with <i>fildev</i> .                                                                                                                                                                                                                                                                                       |
| 36425 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36426 | XSI | The <i>pread()</i> function performs the same action as <i>read()</i> , except that it reads from a given position in the file without changing the file pointer. The first three arguments to <i>pread()</i> are the same as <i>read()</i> with the addition of a fourth argument offset for the desired position inside the file. An attempt to perform a <i>pread()</i> on a file that is incapable of seeking results in an error. |
| 36427 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36428 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36429 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36430 |     | If <i>fildev</i> refers to a socket, <i>read()</i> is equivalent to <i>recv()</i> with no flags set.                                                                                                                                                                                                                                                                                                                                   |
| 36431 |     | <b>RETURN VALUE</b>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 36432 | XSI | Upon successful completion, <i>read()</i> , <i>pread()</i> and <i>readv()</i> shall return a non-negative integer indicating the number of bytes actually read. Otherwise, the functions shall return <i>-1</i> and set <i>errno</i> to indicate the error.                                                                                                                                                                            |
| 36433 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36434 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36435 |     | <b>ERRORS</b>                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 36436 | XSI | The <i>read()</i> , <i>pread()</i> and <i>readv()</i> functions shall fail if:                                                                                                                                                                                                                                                                                                                                                         |
| 36437 |     | [EAGAIN] The <i>O_NONBLOCK</i> flag is set for the file descriptor and the process would be delayed.                                                                                                                                                                                                                                                                                                                                   |
| 36438 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36439 |     | [EBADF] The <i>fildev</i> argument is not a valid file descriptor open for reading.                                                                                                                                                                                                                                                                                                                                                    |
| 36440 | XSR | [EBADMSG] The file is a <i>STREAM</i> file that is set to control-normal mode and the message waiting to be read includes a control part.                                                                                                                                                                                                                                                                                              |
| 36441 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36442 |     | [EINTR] The read operation was terminated due to the receipt of a signal, and no data was transferred.                                                                                                                                                                                                                                                                                                                                 |
| 36443 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36444 | XSR | [EINVAL] The <i>STREAM</i> or multiplexer referenced by <i>fildev</i> is linked (directly or indirectly) downstream from a multiplexer.                                                                                                                                                                                                                                                                                                |
| 36445 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36446 |     | [EIO] The process is a member of a background process attempting to read from its controlling terminal, the process is ignoring or blocking the <i>SIGTTIN</i> signal, or the process group is orphaned. This error may also be generated for implementation-defined reasons.                                                                                                                                                          |
| 36447 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36448 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36449 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36450 | XSI | [EISDIR] The <i>fildev</i> argument refers to a directory and the implementation does not allow the directory to be read using <i>read()</i> , <i>pread()</i> , or <i>readv()</i> . The <i>readdir()</i> function should be used instead.                                                                                                                                                                                              |
| 36451 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36452 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36453 |     | [EOVERFLOW] The file is a regular file, <i>nbyte</i> is greater than 0, the starting position is before the end-of-file, and the starting position is greater than or equal to the offset maximum established in the open file description associated with <i>fildev</i> .                                                                                                                                                             |
| 36454 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36455 |     |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 36456 | MAN | The <i>read()</i> function shall fail if:                                                                                                                                                                                                                                                                                                                                                                                              |

|       |                           |                                                                                              |
|-------|---------------------------|----------------------------------------------------------------------------------------------|
| 36457 | [EAGAIN] or [EWOULDBLOCK] |                                                                                              |
| 36458 |                           | The file descriptor is for a connection-made socket, is marked                               |
| 36459 |                           | O_NONBLOCK, and no data is waiting to be received.                                           |
| 36460 | [ECONNRESET]              | A read was attempted on a connection-mode socket and the connection was                      |
| 36461 |                           | forcibly closed by its peer.                                                                 |
| 36462 | [ENOTCONN]                | A read was attempted on a connection-made socket that is not connected.                      |
| 36463 | [ETIMEDOUT]               | A read was attempted on a connection-mode socket and a transmission                          |
| 36464 |                           | timeout occurred.                                                                            |
| 36465 |                           |                                                                                              |
| 36466 |                           | The <i>readv()</i> function shall fail if:                                                   |
| 36467 | [EINVAL]                  | The sum of the <i>iov_len</i> values in the <i>iov</i> array overflowed an <i>ssize_t</i> .  |
| 36468 | XSI                       | The <i>read()</i> , <i>pread()</i> and <i>readv()</i> functions may fail if:                 |
| 36469 | MAN                       | [EIO] A physical I/O error has occurred.                                                     |
| 36470 | MAN                       | [ENOBUFS] Insufficient resources were available in the system to perform the operation.      |
| 36471 | MAN                       | [ENOMEM] Insufficient memory was available to fulfill the request.                           |
| 36472 |                           | [ENXIO] A request was made of a nonexistent device, or the request was outside the           |
| 36473 |                           | capabilities of the device.                                                                  |
| 36474 |                           | The <i>readv()</i> function may fail if:                                                     |
| 36475 | XSI                       | [EINVAL] The <i>iovcnt</i> argument was less than or equal to 0, or greater than {IOV_MAX}.  |
| 36476 | XSI                       | The <i>pread()</i> function shall fail, and the file pointer shall remain unchanged, if:     |
| 36477 | XSI                       | [EINVAL] The <i>offset</i> argument is invalid. The value is negative.                       |
| 36478 | XSI                       | [EOVERFLOW] The file is a regular file and an attempt was made to read or write at or beyond |
| 36479 |                           | the offset maximum associated with the file.                                                 |
| 36480 | XSI                       | [ENXIO] A request was outside the capabilities of the device.                                |
| 36481 | XSI                       | [ESPIPE] <i>fdes</i> is associated with a pipe or FIFO.                                      |

36482 **EXAMPLES**36483 **Reading Data into a Buffer**

36484 The following example reads data from the file associated with the file descriptor *fd* into the  
 36485 buffer pointed to by *buf*.

```

36486 #include <sys/types.h>
36487 #include <unistd.h>
36488 ...
36489 char buf[20];
36490 size_t nbytes;
36491 ssize_t bytes_read;
36492 int fd;
36493 ...
36494 nbytes = sizeof(buf);
36495 bytes_read = read(fd, buf, nbytes);
36496 ...

```

36497 **Reading Data into an Array**

36498 The following example reads data from the file associated with the file descriptor *fd* into the  
 36499 buffers specified by members of the *iov* array.

```

36500 #include <sys/types.h>
36501 #include <sys/uio.h>
36502 #include <unistd.h>
36503 ...
36504 ssize_t bytes_read;
36505 int fd;
36506 char buf0[20];
36507 char buf1[30];
36508 char buf2[40];
36509 int iovcnt;
36510 struct iovec iov[3];

36511 iov[0].iov_base = buf0;
36512 iov[0].iov_len = sizeof(buf0);
36513 iov[1].iov_base = buf1;
36514 iov[1].iov_len = sizeof(buf1);
36515 iov[2].iov_base = buf2;
36516 iov[2].iov_len = sizeof(buf2);
36517 ...
36518 iovcnt = sizeof(iov) / sizeof(struct iovec);

36519 bytes_read = readv(fd, iov, iovcnt);
36520 ...

```

36521 **APPLICATION USAGE**

36522 None.

36523 **RATIONALE**

36524 This volume of IEEE Std. 1003.1-200x does not specify the value of the file offset after an error is  
 36525 returned; there are too many cases. For programming errors, such as [EBADF], the concept is  
 36526 meaningless since no file is involved. For errors that are detected immediately, such as  
 36527 [EAGAIN], clearly the pointer should not change. After an interrupt or hardware error, however,  
 36528 an updated value would be very useful and is the behavior of many implementations.

36529 Note that a *read()* of zero bytes does not modify *st\_atime*. A *read()* that requests more than zero  
 36530 bytes, but returns zero, shall modify *st\_atime*.

36531 Implementations are allowed, but not required, to perform error checking for *read()* requests of  
 36532 zero bytes.

36533 **Input and Output**

36534 The use of I/O with large byte counts has always presented problems. Ideas such as *lread()* and  
 36535 *lwrite()* (using and returning **longs**) were considered at one time. The current solution is to use  
 36536 abstract types on the ISO C standard function to *read()* and *write()*. The abstract types can be  
 36537 declared so that existing functions work, but can also be declared so that larger types can be  
 36538 represented in future implementations. It is presumed that whatever constraints limit the  
 36539 maximum range of **size\_t** also limit portable I/O requests to the same range. This volume of  
 36540 IEEE Std. 1003.1-200x also limits the range further by requiring that the byte count be limited so  
 36541 that a signed return value remains meaningful. Since the return type is also a (signed) abstract  
 36542 type, the byte count can be defined by the implementation to be larger than an **int** can hold.

36543 The standard developers considered adding atomicity requirements to a pipe or FIFO, but  
36544 recognized that due to the nature of pipes and FIFOs there could be no guarantee of atomicity of  
36545 reads of {PIPE\_BUF} or any other size that would be an aid to applications portability.

36546 This volume of IEEE Std. 1003.1-200x requires that no action be taken when *nbyte* is zero. This is  
36547 not intended to take precedence over detection of errors (such as invalid buffer pointers or file  
36548 descriptors). This is consistent with the rest of this volume of IEEE Std. 1003.1-200x, but the  
36549 phrasing here could be misread to require detection of the zero case before any other errors. A  
36550 value of zero is to be considered a correct value, for which the semantics are a no-op.

36551 I/O is intended to be atomic to ordinary files and pipes and FIFOs. Atomic means that all the  
36552 bytes from a single operation that started out together end up together, without interleaving  
36553 from other I/O operations. It is a known attribute of terminals that this is not honored, and  
36554 terminals are explicitly (and implicitly permanently) excepted, making the behavior unspecified.  
36555 The behavior for other device types is also left unspecified, but the wording is intended to imply  
36556 that future standards might choose to specify atomicity (or not).

36557 There were recommendations to add format parameters to *read()* and *write()* in order to handle  
36558 networked transfers among heterogeneous file system and base hardware types. Such a facility  
36559 may be required for support by the OSI presentation of layer services. However, it was  
36560 determined that this should correspond with similar C-language facilities, and that is beyond the  
36561 scope of this volume of IEEE Std. 1003.1-200x. The concept was suggested to the developers of  
36562 the ISO C standard for their consideration as a possible area for future work.

36563 In 4.3 BSD, a *read()* or *write()* that is interrupted by a signal before transferring any data does not  
36564 by default return an [EINTR] error, but is restarted. In 4.2 BSD, 4.3 BSD, and the Eighth Edition,  
36565 there is an additional function, *select()*, whose purpose is to pause until specified activity (data  
36566 to read, space to write, and so on) is detected on specified file descriptors. It is common in  
36567 applications written for those systems for *select()* to be used before *read()* in situations (such as  
36568 keyboard input) where interruption of I/O due to a signal is desired.

36569 The issue of which files or file types are interruptible is considered an implementation design  
36570 issue. This is often affected primarily by hardware and reliability issues.

36571 There are no references to actions taken following an “unrecoverable error”. It is considered  
36572 beyond the scope of this volume of IEEE Std. 1003.1-200x to describe what happens in the case of  
36573 hardware errors.

36574 Previous versions of IEEE Std. 1003.1-200x allowed two very different behaviors with regard to  
36575 the handling of interrupts. In order to minimize the resulting confusion, it was decided that  
36576 IEEE Std. 1003.1-200x should support only one of these behaviors. Historical practice on AT&T-  
36577 derived systems was to have *read()* and *write()* return  $-1$  and set *errno* to [EINTR] when  
36578 interrupted after some, but not all, of the data requested had been transferred. However, the U.S.  
36579 Department of Commerce FIPS 151-1 and FIPS 151-2 require the historical BSD behavior, in  
36580 which *read()* and *write()* return the number of bytes actually transferred before the interrupt. If  
36581  $-1$  is returned when any data is transferred, it is difficult to recover from the error on a seekable  
36582 device and impossible on a non-seekable device. Most new implementations support this  
36583 behavior. The behavior required by IEEE Std. 1003.1-200x is to return the number of bytes  
36584 transferred.

36585 IEEE Std. 1003.1-200x does not specify when an implementation that buffers *read()*s actually  
36586 moves the data into the user-supplied buffer, so an implementation may choose to do this at the  
36587 latest possible moment. Therefore, an interrupt arriving earlier may not cause *read()* to return a  
36588 partial byte count, but rather to return  $-1$  and set *errno* to [EINTR].

36589 Consideration was also given to combining the two previous options, and setting *errno* to  
36590 [EINTR] while returning a short count. However, not only is there no existing practice that

36591 implements this, it is also contradictory to the idea that when *errno* is set, the function  
36592 responsible shall return  $-1$ .

#### 36593 FUTURE DIRECTIONS

36594 None.

#### 36595 SEE ALSO

36596 *fcntl()*, *ioctl()*, *lseek()*, *open()*, *pipe()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
36597 *<stropts.h>*, *<sys/uio.h>*, *<unistd.h>*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
36598 Chapter 11, General Terminal Interface

#### 36599 CHANGE HISTORY

36600 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 36601 Issue 4

36602 The *<unistd.h>* header is added to the SYNOPSIS section.

36603 The DESCRIPTION is rearranged for clarity and to align more closely with the ISO POSIX-1  
36604 standard. No functional changes are made other than as noted elsewhere in this CHANGE  
36605 HISTORY section.

36606 In the ERRORS section in previous issues, generation of the [EIO] error depended on whether or  
36607 not an implementation supported Job Control. This functionality is now defined as mandatory.

36608 The [ENXIO] error is marked as an extension.

36609 The APPLICATION USAGE section is removed.

36610 The description of [EINTR] is amended.

36611 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 36612 • The type of the argument *buf* is changed from **char\*** to **void\***, and the type of the argument  
36613 *nbyte* is changed from **unsigned** to **size\_t**.

- 36614 • The DESCRIPTION now states that the result is implementation-defined if *nbyte* is greater  
36615 than {SSIZE\_MAX}. This limit was defined by the constant {INT\_MAX} in Issue 3.

36616 The following change is incorporated for alignment with the FIPS requirements:

- 36617 • The last paragraph of the DESCRIPTION now states that if *read()* is interrupted by a signal  
36618 after it has successfully read some data, it returns the number of bytes read. In Issue 3, it was  
36619 optional whether *read()* returned the number of bytes read, or whether it returned  $-1$  with  
36620 *errno* set to [EINTR].

#### 36621 Issue 4, Version 2

36622 The following changes are incorporated for X/OPEN UNIX conformance:

- 36623 • The *readv()* function is added to the SYNOPSIS.

- 36624 • The DESCRIPTION is updated to describe the reading of data from STREAMS files. An  
36625 operational description of the *readv()* function is also added.

- 36626 • References to the *readv()* function are added to the RETURN VALUE and ERRORS sections  
36627 in appropriate places.

- 36628 • The ERRORS section has been restructured to describe errors that apply generally (that is, to  
36629 both *read()* and *readv()*), and to describe those that apply to *readv()* specifically. The  
36630 [EBADMSG], [EINVAL], and [EISDIR] errors are also added.

36631 **Issue 5**

36632 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
36633 Threads Extension.

36634 Large File Summit extensions are added.

36635 The *pread()* function is added.

36636 **Issue 6**

36637 The DESCRIPTION and ERRORS sections are updated so that references to STREAMS are  
36638 marked as part of the XSI STREAMS Option Group.

36639 The following new requirements on POSIX implementations derive from alignment with the  
36640 Single UNIX Specification:

36641 • The DESCRIPTION now states that if *read()* is interrupted by a signal after it has successfully  
36642 read some data, it returns the number of bytes read. In Issue 3, it was optional whether *read()*  
36643 returned the number of bytes read, or whether it returned  $-1$  with *errno* set to [EINTR]. This  
36644 is a FIPS requirement.

36645 • In the DESCRIPTION, text is added to indicate that for regular files, no data transfer occurs  
36646 past the offset maximum established in the open file description associated with *files*. This  
36647 change is to support large files.

36648 • The [EOVERFLOW] mandatory error condition is added.

36649 • The [ENXIO] optional error condition is added.

36650 Text referring to sockets is added to the DESCRIPTION.

36651 The following changes were made to align with the IEEE P1003.1a draft standard:

36652 • The effect of reading zero bytes is clarified.

36653 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that  
36654 *read()* results are unspecified for typed memory objects.

36655 New RATIONALE is added to explain the atomicity requirements for input and output  
36656 operations.

36657 The following error conditions are added for operations on sockets: [EAGAIN],  
36658 [ECONNRESET], [ENOTCONN], and [ETIMEDOUT].

36659 The [EIO] error is changed to “may fail”.

36660 The following error conditions are added for operations on sockets: [ENOBUFFS] and  
36661 [ENOMEM].

## 36662 NAME

36663 readdir, readdir\_r — read directory

## 36664 SYNOPSIS

36665 #include &lt;dirent.h&gt;

36666 struct dirent \*readdir(DIR \*dirp);

36667 TSF int readdir\_r(DIR \*restrict dirp, struct dirent \*restrict entry,

36668 struct dirent \*\*restrict result);

36669

## 36670 DESCRIPTION

36671 The type **DIR**, which is defined in the header <dirent.h>, represents a *directory stream*, which is  
 36672 an ordered sequence of all the directory entries in a particular directory. Directory entries  
 36673 represent files; files may be removed from a directory or added to a directory asynchronously to  
 36674 the operation of *readdir()*.

36675 The *readdir()* function shall return a pointer to a structure representing the directory entry at the  
 36676 current position in the directory stream specified by the argument *dirp*, and position the  
 36677 directory stream at the next entry. It shall return a null pointer upon reaching the end of the  
 36678 directory stream. The structure **dirent** defined by the <dirent.h> header describes a directory  
 36679 entry.

36680 The *readdir()* function shall not return directory entries containing empty names. If entries for  
 36681 dot or dot-dot exist, one entry shall be returned for dot and one entry shall be returned for dot-  
 36682 dot; otherwise, they shall not be returned.

36683 The pointer returned by *readdir()* points to data which may be overwritten by another call to  
 36684 *readdir()* on the same directory stream. This data is not overwritten by another call to *readdir()*  
 36685 on a different directory stream.

36686 If a file is removed from or added to the directory after the most recent call to *opendir()* or  
 36687 *rewinddir()*, whether a subsequent call to *readdir()* returns an entry for that file is unspecified.

36688 The *readdir()* function may buffer several directory entries per actual read operation; *readdir()*  
 36689 shall mark for update the *st\_atime* field of the directory each time the directory is actually read.

36690 After a call to *fork()*, either the parent or child (but not both) may continue processing the  
 36691 XSI directory stream using *readdir()*, *rewinddir()*, or *seekdir()*. If both the parent and child processes  
 36692 use these functions, the result is undefined.

36693 If the entry names a symbolic link, the value of the *d\_ino* member is unspecified.

36694 The *readdir()* function need not be reentrant. A function that is not required to be reentrant is not  
 36695 required to be thread-safe.

36696 TSF The *readdir\_r()* function initializes the **dirent** structure referenced by *entry* to represent the  
 36697 directory entry at the current position in the directory stream referred to by *dirp*, store a pointer  
 36698 to this structure at the location referenced by *result*, and position the directory stream at the next  
 36699 entry.

36700 The storage pointed to by *entry* shall be large enough for a **dirent** with an array of **char** *d\_name*  
 36701 members containing at least {NAME\_MAX} plus one elements.

36702 Upon successful return, the pointer returned at *\*result* shall have the same value as the argument  
 36703 *entry*. Upon reaching the end of the directory stream, this pointer shall have the value NULL.

36704 The *readdir\_r()* function shall not return directory entries containing empty names.

36705 If a file is removed from or added to the directory after the most recent call to *opendir()* or  
 36706 *rewinddir()*, whether a subsequent call to *readdir\_r()* returns an entry for that file is unspecified.

36707 The *readdir\_r()* function may buffer several directory entries per actual read operation; the  
 36708 *readdir\_r()* function shall mark for update the *st\_atime* field of the directory each time the  
 36709 directory is actually read.

36710 Applications wishing to check for error situations should set *errno* to 0 before calling *readdir()*. If  
 36711 *errno* is set to non-zero on return, an error occurred.

#### 36712 RETURN VALUE

36713 Upon successful completion, *readdir()* shall return a pointer to an object of type **struct dirent**.  
 36714 When an error is encountered, a null pointer shall be returned and *errno* shall be set to indicate  
 36715 the error. When the end of the directory is encountered, a null pointer shall be returned and *errno*  
 36716 is not changed.

36717 TSF If successful, the *readdir\_r()* function shall return zero; otherwise, an error number shall be  
 36718 returned to indicate the error.

#### 36719 ERRORS

36720 The *readdir()* function shall fail if:

36721 [EOVERFLOW] One of the values in the structure to be returned cannot be represented  
 36722 correctly.

36723 The *readdir()* function may fail if:

36724 [EBADF] The *dirp* argument does not refer to an open directory stream.

36725 [ENOENT] The current position of the directory stream is invalid.

36726 The *readdir\_r()* function may fail if:

36727 [EBADF] The *dirp* argument does not refer to an open directory stream.

#### 36728 EXAMPLES

36729 The following sample code searches the current directory for the entry *name*:

```
36730 dirp = opendir(".");
36731 while (dirp) {
36732 errno = 0;
36733 if ((dp = readdir(dirp)) != NULL) {
36734 if (strcmp(dp->d_name, name) == 0) {
36735 closedir(dirp);
36736 return FOUND;
36737 }
36738 } else {
36739 if (errno == 0) {
36740 closedir(dirp);
36741 return NOT_FOUND;
36742 }
36743 closedir(dirp);
36744 return READ_ERROR;
36745 }
36746 }
36747 return OPEN_ERROR;
```

36748 **APPLICATION USAGE**

36749 The *readdir()* function should be used in conjunction with *opendir()*, *closedir()*, and *rewinddir()* to  
 36750 examine the contents of the directory.

36751 The *readdir\_r()* function is thread-safe and shall return values in a user-supplied buffer instead  
 36752 of possibly using a static data area that may be overwritten by each call.

36753 **RATIONALE**

36754 The returned value of *readdir()* merely *represents* a directory entry. No equivalence should be  
 36755 inferred.

36756 Historical implementations of *readdir()* obtain multiple directory entries on a single read  
 36757 operation, which permits subsequent *readdir()* operations to operate from the buffered  
 36758 information. Any wording that required each successful *readdir()* operation to mark the  
 36759 directory *st\_atime* field for update would militate against the historical performance-oriented  
 36760 implementations.

36761 Since *readdir()* returns NULL when it detects an error and when the end of the directory is  
 36762 encountered, an application that needs to tell the difference must set *errno* to zero before the call  
 36763 and check it if NULL is returned. Because the function must not change *errno* in the second case  
 36764 and must set it to a non-zero value in the first case, a zero *errno* after a call returning NULL  
 36765 indicates end of directory; otherwise, an error.

36766 Routines to deal with this problem more directly were proposed:

```
36767 int derror (dirp)
36768 DIR *dirp;
```

```
36769 void clearerr (dirp)
36770 DIR *dirp;
```

36771 The first would indicate whether an error had occurred, and the second would clear the error  
 36772 indication. The simpler method involving *errno* was adopted instead by requiring that *readdir()*  
 36773 not change *errno* when end-of-directory is encountered.

36774 An error or signal indicating that a directory has changed while open was considered but  
 36775 rejected.

36776 The thread-safe version of the directory reading function shall return values in a user-supplied  
 36777 buffer instead of possibly using a static data area that may be overwritten by each call. Either the  
 36778 {NAME\_MAX} compile-time constant or the corresponding *pathconf()* option can be used to  
 36779 determine the maximum sizes of returned path names.

36780 **FUTURE DIRECTIONS**

36781 None.

36782 **SEE ALSO**

36783 *closedir()*, *lstat()*, *opendir()*, *rewinddir()*, *symlink()*, the Base Definitions volume of  
 36784 IEEE Std. 1003.1-200x, <*dirent.h*>, <*sys/types.h*>

36785 **CHANGE HISTORY**

36786 First released in Issue 2.

36787 **Issue 4**

36788 The <*sys/types.h*> header is now marked as optional (OH); this header need not be included on  
 36789 XSI-conformant systems.

36790 In the DESCRIPTION, the fact that XSI-conformant systems return entries for dot and dot-dot is  
 36791 marked as an extension. This functionality is not specified in the ISO POSIX-1 standard.

- 36792 There is some rewording of the DESCRIPTION and RETURN VALUE sections. No functional  
36793 changes are made other than as noted elsewhere in this CHANGE HISTORY section.
- 36794 The following change is incorporated for alignment with the ISO POSIX-1 standard:
- 36795
- The last paragraph of the DESCRIPTION describing a restriction after *fork()* is added.
- 36796 **Issue 4, Version 2**
- 36797 The following changes are incorporated for X/OPEN UNIX conformance:
- 36798
- A statement is added to the DESCRIPTION indicating the disposition of certain fields in  
36799 **struct dirent** when an entry refers to a symbolic link.
- 36800
- The [ENOENT] optional error condition is added.
- 36801 **Issue 5**
- 36802 Large File Summit extensions are added.
- 36803 The *readdir\_r()* function is included for alignment with the POSIX Threads Extension.
- 36804 A note indicating that the *readdir()* function need not be reentrant is added to the  
36805 DESCRIPTION.
- 36806 **Issue 6**
- 36807 The *readdir\_r()* function is marked as part of the Thread-Safe Functions option.
- 36808 The Open Group corrigenda item U026/7 has been applied, correcting the prototype for  
36809 *readdir\_r()*.
- 36810 The Open Group corrigenda item U026/8 has been applied, clarifying the wording of the  
36811 successful return for the *readdir\_r()* function.
- 36812 The following new requirements on POSIX implementations derive from alignment with the  
36813 Single UNIX Specification:
- 36814
- The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
36815 required for conforming implementations of previous POSIX specifications, it was not  
36816 required for UNIX applications.
- 36817
- A statement is added to the DESCRIPTION indicating the disposition of certain fields in  
36818 **struct dirent** when an entry refers to a symbolic link.
- 36819
- The [EOVERFLOW] mandatory error condition is added. This change is to support large  
36820 files.
- 36821
- The [ENOENT] optional error condition is added.
- 36822 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
36823 its avoidance of possibly using a static data area.
- 36824 The **restrict** keyword is added to the *readdir\_r()* prototype for alignment with the  
36825 ISO/IEC 9899:1999 standard.

36826 **NAME**

36827 readlink — read the contents of a symbolic link

36828 **SYNOPSIS**

36829 #include &lt;unistd.h&gt;

36830 ssize\_t readlink(const char \*restrict path, char \*restrict buf,  
36831 size\_t bufsize);36832 **DESCRIPTION**36833 The *readlink()* function shall place the contents of the symbolic link referred to by *path* in the  
36834 buffer *buf* which has size *bufsize*. If the number of bytes in the symbolic link is less than *bufsize*,  
36835 the contents of the remainder of *buf* are unspecified. If the *buf* argument is not large enough to  
36836 contain the link content, the first *bufsize* bytes shall be placed in *buf*.36837 If the value of *bufsize* is greater than {SSIZE\_MAX}, the result is implementation-defined.36838 **RETURN VALUE**36839 Upon successful completion, *readlink()* shall return the count of bytes placed in the buffer.  
36840 Otherwise, it shall return a value of -1, leave the buffer unchanged, and set *errno* to indicate the  
36841 error.36842 **ERRORS**36843 The *readlink()* function shall fail if:36844 [EACCES] Search permission is denied for a component of the path prefix of *path*.36845 [EINVAL] The *path* argument names a file that is not a symbolic link.

36846 [EIO] An I/O error occurred while reading from the file system.

36847 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
36848 argument.

36849 [ENAMETOOLONG]

36850 The length of the *path* argument exceeds {PATH\_MAX} or a path name  
36851 component is longer than {NAME\_MAX}.36852 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

36853 [ENOTDIR] A component of the path prefix is not a directory.

36854 The *readlink()* function may fail if:

36855 [EACCES] Read permission is denied for the directory.

36856 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
36857 resolution of the *path* argument.

36858 [ENAMETOOLONG]

36859 As a result of encountering a symbolic link in resolution of the *path* argument,  
36860 the length of the substituted path name string exceeded {PATH\_MAX}.

36861 **EXAMPLES**36862 **Reading the Name of a Symbolic Link**

36863 The following example shows how to read the name of a symbolic link named `/modules/pass1`.

```
36864 #include <unistd.h>
36865 char buf[1024];
36866 int len;
36867 ...
36868 if ((len = readlink("/modules/pass1", buf, sizeof(buf)-1)) != -1);
36869 buf[len] = '\0';
```

36870 **APPLICATION USAGE**

36871 Portable applications should not assume that the returned contents of the symbolic link are  
36872 null-terminated.

36873 **RATIONALE**

36874 Since IEEE Std. 1003.1-200x does not require any association of file times with symbolic links,  
36875 there is no requirement that file times be updated by `readlink()`. The type associated with `bufsiz`  
36876 is a `size_t` in order to be consistent with both the ISO C standard and the definition of `read()`.  
36877 The behavior specified for `readlink()` when `bufsiz` is zero represents historical practice. For this  
36878 case, the standard developers considered a change whereby `readlink()` would return the number  
36879 of non-null bytes contained in the symbolic link with the buffer `buf` remaining unchanged;  
36880 however, since the `stat` structure member `st_size` value can be used to determine the size of  
36881 buffer necessary to contain the contents of the symbolic link as returned by `readlink()`, this  
36882 proposal was rejected, and the historical practice retained.

36883 **FUTURE DIRECTIONS**

36884 None.

36885 **SEE ALSO**

36886 `lstat()`, `stat()`, `symlink()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<unistd.h>`

36887 **CHANGE HISTORY**

36888 First released in Issue 4, Version 2.

36889 **Issue 5**

36890 Moved from X/OPEN UNIX extension to BASE.

36891 **Issue 6**

36892 The return type is changed to `ssize_t`, to align with the IEEE P1003.1a draft standard.

36893 The following new requirements on POSIX implementations derive from alignment with the  
36894 Single UNIX Specification:

- 36895 • This function is made mandatory.
- 36896 • In this function it is possible for the return value to exceed the range of the type `ssize_t` (since  
36897 `size_t` has a larger range of positive values than `ssize_t`). A sentence restricting the size of  
36898 the `size_t` object is added to the description to resolve this conflict.

36899 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 36900 • The `[ENAMETOOLONG]` error is restored as an error dependent on `_POSIX_NO_TRUNC`.  
36901 This is since behavior may vary from one file system to another.
- 36902 • The **FUTURE DIRECTIONS** section is changed to None.

36903 The following changes were made to align with the IEEE P1003.1a draft standard:  
36904     • The [ELOOP] optional error condition is added.  
36905 The **restrict** keyword is added to the *readlink()* prototype for alignment with the  
36906 ISO/IEC 9899:1999 standard.

36907 **NAME**

36908           readv — vectored read from file

36909 **SYNOPSIS**

36910 xSI       #include <sys/uio.h>

36911           ssize\_t readv(int *fd*, const struct iovec \**iov*, int *iovcnt*);

36912

36913 **DESCRIPTION**

36914           Refer to *read()*.

36915 **NAME**

36916           realloc — memory reallocator

36917 **SYNOPSIS**

36918           #include &lt;stdlib.h&gt;

36919           void \*realloc(void \*ptr, size\_t size);

36920 **DESCRIPTION**

36921 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
36922 conflict between the requirements described here and the ISO C standard is unintentional. This  
36923 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

36924       The *realloc()* function shall change the size of the memory object pointed to by *ptr* to the size  
36925 specified by *size*. The contents of the object shall remain unchanged up to the lesser of the new  
36926 and old sizes. If the new size of the memory object would require movement of the object, the  
36927 space for the previous instantiation of the object is freed. If the new size is larger, the contents of  
36928 the newly allocated portion of the object are unspecified. If *size* is 0 and *ptr* is not a null pointer,  
36929 the object pointed to is freed. If the space cannot be allocated, the object remains unchanged.

36930       If *ptr* is a null pointer, *realloc()* shall behave like *malloc()* for the specified size.

36931       If *ptr* does not match a pointer returned earlier by *calloc()*, *malloc()*, or *realloc()* or if the space has  
36932 previously been deallocated by a call to *free()* or *realloc()*, the behavior is undefined.

36933       The order and contiguity of storage allocated by successive calls to *realloc()* is unspecified. The  
36934 pointer returned if the allocation succeeds is suitably aligned so that it may be assigned to a  
36935 pointer to any type of object and then used to access such an object in the space allocated (until  
36936 the space is explicitly freed or reallocated). Each such allocation shall yield a pointer to an object  
36937 disjoint from any other object. The pointer returned points to the start (lowest byte address) of  
36938 the allocated space. If the space cannot be allocated, a null pointer shall be returned.

36939 **RETURN VALUE**

36940       Upon successful completion with a *size* not equal to 0, *realloc()* shall return a pointer to the  
36941 (possibly moved) allocated space. If *size* is 0, either a null pointer or a unique pointer that can be  
36942 successfully passed to *free()* shall be returned. If there is not enough available memory, *realloc()*  
36943 **CX** shall return a null pointer and set *errno* to [ENOMEM].

36944 **ERRORS**

36945       The *realloc()* function shall fail if:

36946 **CX**       [ENOMEM]       Insufficient memory is available.

36947 **EXAMPLES**

36948       None.

36949 **APPLICATION USAGE**

36950       None.

36951 **RATIONALE**

36952       None.

36953 **FUTURE DIRECTIONS**

36954       None.

36955 **SEE ALSO**

36956       *calloc()*, *free()*, *malloc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

36957 **CHANGE HISTORY**

36958 First released in Issue 1. Derived from Issue 1 of the SVID.

36959 **Issue 4**36960 The setting of *errno* and the [ENOMEM] error are marked as extensions.

36961 The APPLICATION USAGE section is removed.

36962 The following changes are incorporated for alignment with the ISO C standard:

- 36963 • The DESCRIPTION is updated to indicate as follows:
  - 36964 — The order and contiguity of storage allocated by successive calls to this function is
  - 36965 unspecified.
  - 36966 — Each allocation yields a pointer to an object disjoint from any other object.
  - 36967 — The returned pointer points to the lowest byte address of the allocation.
- 36968 • The RETURN VALUE section is updated to indicate what is returned if *size* is 0.

36969 **Issue 6**

36970 Extensions beyond the ISO C standard are now marked.

36971 The following new requirements on POSIX implementations derive from alignment with the  
36972 Single UNIX Specification:

- 36973 • In the RETURN VALUE section, if there is not enough available memory, the setting of *errno*
- 36974 to [ENOMEM] is added.
- 36975 • The [ENOMEM] error condition is added.

36976 **NAME**

36977           realpath — resolve a path name

36978 **SYNOPSIS**

36979 XSI       #include &lt;stdlib.h&gt;

36980           char \*realpath(const char \*restrict *file\_name*,  
36981                        char \*restrict *resolved\_name*);

36982

36983 **DESCRIPTION**

36984       The *realpath()* function derives, from the path name pointed to by *file\_name*, an absolute path name that names the same file, whose resolution does not involve '.', '..', or symbolic links.

36985       The generated path name is stored as a null-terminated string, up to a maximum of {PATH\_MAX} bytes, in the buffer pointed to by *resolved\_name*.

36988 **RETURN VALUE**

36989       Upon successful completion, *realpath()* shall return a pointer to the resolved name. Otherwise, *realpath()* shall return a null pointer and set *errno* to indicate the error, and the contents of the buffer pointed to by *resolved\_name* are undefined.

36992 **ERRORS**36993       The *realpath()* function shall fail if:

- |       |                |                                                                                                                    |
|-------|----------------|--------------------------------------------------------------------------------------------------------------------|
| 36994 | [EACCES]       | Read or search permission was denied for a component of <i>file_name</i> .                                         |
| 36995 | [EINVAL]       | Either the <i>file_name</i> or <i>resolved_name</i> argument is a null pointer.                                    |
| 36996 | [EIO]          | An error occurred while reading from the file system.                                                              |
| 36997 | [ELOOP]        | A loop exists in symbolic links encountered during resolution of the <i>path</i> argument.                         |
| 36998 |                |                                                                                                                    |
| 36999 | [ENAMETOOLONG] | The length of the <i>file_name</i> argument exceeds {PATH_MAX} or a path name component is longer than {NAME_MAX}. |
| 37000 |                |                                                                                                                    |
| 37001 |                |                                                                                                                    |
| 37002 | [ENOENT]       | A component of <i>file_name</i> does not name an existing file or <i>file_name</i> points to an empty string.      |
| 37003 |                |                                                                                                                    |
| 37004 | [ENOTDIR]      | A component of the path prefix is not a directory.                                                                 |
| 37005 |                | The <i>realpath()</i> function may fail if:                                                                        |
| 37006 | [ELOOP]        | More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.             |
| 37007 |                |                                                                                                                    |
| 37008 | [ENAMETOOLONG] | Path name resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.           |
| 37009 |                |                                                                                                                    |
| 37010 |                |                                                                                                                    |
| 37011 | [ENOMEM]       | Insufficient storage space is available.                                                                           |

37012 **EXAMPLES**37013 **Generating an Absolute Path Name**

37014 The following example generates an absolute path name for the file identified by the *symlinkpath*  
37015 argument. The generated path name is stored in the *actualpath* array.

```
37016 #include <stdlib.h>
37017 ...
37018 char *symlinkpath = "/tmp/symlink/file";
37019 char actualpath [PATH_MAX+1];
37020 char *ptr;

37021 ptr = realpath(symlinkpath, actualpath);
```

37022 **APPLICATION USAGE**

37023 None.

37024 **RATIONALE**

37025 None.

37026 **FUTURE DIRECTIONS**

37027 None.

37028 **SEE ALSO**

37029 *getcwd()*, *sysconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdlib.h**>

37030 **CHANGE HISTORY**

37031 First released in Issue 4, Version 2.

37032 **Issue 5**

37033 Moved from X/OPEN UNIX extension to BASE.

37034 **Issue 6**

37035 The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`. This  
37036 is since behavior may vary from one file system to another.

37037 The **restrict** keyword is added to the *realpath()* prototype for alignment with the  
37038 ISO/IEC 9899:1999 standard.

37039 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
37040 [ELOOP] error condition is added.

37041 **NAME**

37042           recv — receive a message from a connected socket

37043 **SYNOPSIS**

37044           #include &lt;sys/socket.h&gt;

37045           ssize\_t recv(int *socket*, void \**buffer*, size\_t *length*, int *flags*);37046 **DESCRIPTION**

37047           The *recv()* function receives a message from a connection-mode or connectionless-mode socket.  
 37048           It is normally used with connected sockets because it does not permit the application to retrieve  
 37049           the source address of received data.

37050           The *recv()* function takes the following arguments:37051           *socket*           Specifies the socket file descriptor.37052           *buffer*           Points to a buffer where the message should be stored.37053           *length*           Specifies the length in bytes of the buffer pointed to by the *buffer* argument.

37054           *flags*            Specifies the type of message reception. Values of this argument are formed by  
 37055           logically OR'ing zero or more of the following values:

37056                   MSG\_PEEK       Peeks at an incoming message. The data is treated as unread and  
 37057                   the next *recv()* or similar function shall still return this data.

37058                   MSG\_OOB       Requests out-of-band data. The significance and semantics of  
 37059                   out-of-band data are protocol-specific.

37060                   MSG\_WAITALL Requests that the function block until the full amount of data  
 37061                   requested can be returned. The function may return a smaller  
 37062                   amount of data if a signal is caught, if the connection is  
 37063                   terminated, if MSG\_PEEK was specified, or if an error is pending  
 37064                   for the socket.

37065           The *recv()* function shall return the length of the message written to the buffer pointed to by the  
 37066           *buffer* argument. For message-based sockets, such as SOCK\_DGRAM and SOCK\_SEQPACKET,  
 37067           the entire message shall be read in a single operation. If a message is too long to fit in the  
 37068           supplied buffer, and MSG\_PEEK is not set in the *flags* argument, the excess bytes shall be  
 37069           discarded. For stream-based sockets, such as SOCK\_STREAM, message boundaries shall be  
 37070           ignored. In this case, data is returned to the user as soon as it becomes available, and no data  
 37071           shall be discarded.

37072           If the MSG\_WAITALL flag is not set, data shall be returned only up to the end of the first  
 37073           message.

37074           If no messages are available at the socket and O\_NONBLOCK is not set on the socket's file  
 37075           descriptor, *recv()* shall block until a message arrives. If no messages are available at the socket  
 37076           and O\_NONBLOCK is set on the socket's file descriptor, *recv()* shall fail and set *errno* to  
 37077           [EAGAIN] or [EWOULDBLOCK].

37078 **RETURN VALUE**

37079           Upon successful completion, *recv()* shall return the length of the message in bytes. If no  
 37080           messages are available to be received and the peer has performed an orderly shutdown, *recv()*  
 37081           shall return 0. Otherwise, -1 shall be returned and *errno* set to indicate the error.

37082 **ERRORS**

- 37083       The *recv()* function shall fail if:
- 37084       [EAGAIN] or [EWOULDBLOCK]
- 37085               The socket's file descriptor is marked O\_NONBLOCK and no data is waiting  
37086               to be received; or MSG\_OOB is set and no out-of-band data is available and  
37087               either the socket's file descriptor is marked O\_NONBLOCK or the socket does  
37088               not support blocking to await out-of-band data.
- 37089       [EBADF]       The *socket* argument is not a valid file descriptor.
- 37090       [ECONNRESET] A connection was forcibly closed by a peer.
- 37091       [EINTR]       The *recv()* function was interrupted by a signal that was caught, before any  
37092               data was available.
- 37093       [EINVAL]       The MSG\_OOB flag is set and no out-of-band data is available.
- 37094       [ENOTCONN]    A receive is attempted on a connection-mode socket that is not connected.
- 37095       [ENOTSOCK]    The *socket* argument does not refer to a socket.
- 37096       [EOPNOTSUPP] The specified flags are not supported for this socket type or protocol.
- 37097       [ETIMEDOUT]    The connection timed out during connection establishment, or due to a  
37098               transmission timeout on active connection.
- 37099       The *recv()* function may fail if:
- 37100       [EIO]         An I/O error occurred while reading from or writing to the file system.
- 37101       [ENOBUFS]     Insufficient resources were available in the system to perform the operation.
- 37102       [ENOMEM]     Insufficient memory was available to fulfill the request.

37103 **EXAMPLES**

37104       None.

37105 **APPLICATION USAGE**

37106       The *recv()* function is identical to *recvfrom()* with a zero *address\_len* argument, and to *read()* if no  
37107       flags are used.

37108       The *select()* and *poll()* functions can be used to determine when data is available to be received.

37109 **RATIONALE**

37110       None.

37111 **FUTURE DIRECTIONS**

37112       None.

37113 **SEE ALSO**

37114       *poll()*, *read()*, *recvmsg()*, *recvfrom()*, *select()*, *send()*, *sendmsg()*, *sendto()*, *shutdown()*, *socket()*,  
37115       *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/socket.h>

37116 **CHANGE HISTORY**

37117       First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

37118 **NAME**

37119           recvfrom — receive a message from a socket

37120 **SYNOPSIS**

```
37121 #include <sys/socket.h>
37122 ssize_t recvfrom(int socket, void *restrict buffer, size_t length,
37123 int flags, struct sockaddr *restrict address,
37124 socklen_t *restrict address_len);
```

37125 **DESCRIPTION**

37126       The *recvfrom()* function receives a message from a connection-mode or connectionless-mode  
 37127       socket. It is normally used with connectionless-mode sockets because it permits the application  
 37128       to retrieve the source address of received data.

37129       The *recvfrom()* function takes the following arguments:

|       |                    |                                                                                                                                                                                                                                                                                                             |
|-------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 37130 | <i>socket</i>      | Specifies the socket file descriptor.                                                                                                                                                                                                                                                                       |
| 37131 | <i>buffer</i>      | Points to the buffer where the message should be stored.                                                                                                                                                                                                                                                    |
| 37132 | <i>length</i>      | Specifies the length in bytes of the buffer pointed to by the <i>buffer</i> argument.                                                                                                                                                                                                                       |
| 37133 | <i>flags</i>       | Specifies the type of message reception. Values of this argument are formed<br>37134 by logically OR'ing zero or more of the following values:                                                                                                                                                              |
| 37135 | MSG_PEEK           | Peeks at an incoming message. The data is treated as unread<br>37136 and the next <i>recvfrom()</i> or similar function shall still return<br>37137 this data.                                                                                                                                              |
| 37138 | MSG_OOB            | Requests out-of-band data. The significance and semantics<br>37139 of out-of-band data are protocol-specific.                                                                                                                                                                                               |
| 37140 | MSG_WAITALL        | Requests that the function block until the full amount of<br>37141 data requested can be returned. The function may return a<br>37142 smaller amount of data if a signal is caught, if the<br>37143 connection is terminated, if MSG_PEEK was specified, or if<br>37144 an error is pending for the socket. |
| 37145 | <i>address</i>     | A null pointer, or points to a <b>sockaddr</b> structure in which the sending address<br>37146 is to be stored. The length and format of the address depend on the address<br>37147 family of the socket.                                                                                                   |
| 37148 | <i>address_len</i> | Specifies the length of the <b>sockaddr</b> structure pointed to by the <i>address</i><br>37149 argument.                                                                                                                                                                                                   |

37150       The *recvfrom()* function shall return the length of the message written to the buffer pointed to by  
 37151       the *buffer* argument. For message-based sockets, such as SOCK\_DGRAM and  
 37152       SOCK\_SEQPACKET, the entire message shall be read in a single operation. If a message is too  
 37153       long to fit in the supplied buffer, and MSG\_PEEK is not set in the *flags* argument, the excess  
 37154       bytes shall be discarded. For stream-based sockets, such as SOCK\_STREAM, message  
 37155       boundaries shall be ignored. In this case, data is returned to the user as soon as it becomes  
 37156       available, and no data shall be discarded.

37157       If the MSG\_WAITALL flag is not set, data shall be returned only up to the end of the first  
 37158       message.

37159       Not all protocols provide the source address for messages. If the *address* argument is not a null  
 37160       pointer and the protocol provides the source address of messages, the source address of the  
 37161       received message is stored in the **sockaddr** structure pointed to by the *address* argument, and the

- 37162 length of this address is stored in the object pointed to by the *address\_len* argument.
- 37163 If the actual length of the address is greater than the length of the supplied **sockaddr** structure,  
37164 the stored address shall be truncated.
- 37165 If the *address* argument is not a null pointer and the protocol does not provide the source address  
37166 of messages, the value stored in the object pointed to by *address* is unspecified.
- 37167 If no messages are available at the socket and O\_NONBLOCK is not set on the socket's file  
37168 descriptor, *recvfrom()* blocks until a message arrives. If no messages are available at the socket  
37169 and O\_NONBLOCK is set on the socket's file descriptor, *recvfrom()* shall fail and set *errno* to  
37170 [EAGAIN] or [EWOULDBLOCK].
- 37171 **RETURN VALUE**
- 37172 Upon successful completion, *recvfrom()* shall return the length of the message in bytes. If no  
37173 messages are available to be received and the peer has performed an orderly shutdown,  
37174 *recvfrom()* shall return 0. Otherwise, the function shall return -1 and set *errno* to indicate the  
37175 error.
- 37176 **ERRORS**
- 37177 The *recvfrom()* function shall fail if:
- 37178 [EAGAIN] or [EWOULDBLOCK] The socket's file descriptor is marked O\_NONBLOCK and no data is waiting  
37179 to be received; or MSG\_OOB is set and no out-of-band data is available and  
37180 either the socket's file descriptor is marked O\_NONBLOCK or the socket does  
37181 not support blocking to await out-of-band data.  
37182
- 37183 [EBADF] The *socket* argument is not a valid file descriptor.
- 37184 [ECONNRESET] A connection was forcibly closed by a peer.
- 37185 [EINTR] A signal interrupted *recvfrom()* before any data was available.
- 37186 [EINVAL] The MSG\_OOB flag is set and no out-of-band data is available.
- 37187 [ENOTCONN] A receive is attempted on a connection-mode socket that is not connected.
- 37188 [ENOTSOCK] The *socket* argument does not refer to a socket.
- 37189 [EOPNOTSUPP] The specified flags are not supported for this socket type.
- 37190 [ETIMEDOUT] The connection timed out during connection establishment, or due to a  
37191 transmission timeout on active connection.
- 37192 The *recvfrom()* function may fail if:
- 37193 [EIO] An I/O error occurred while reading from or writing to the file system.
- 37194 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 37195 [ENOMEM] Insufficient memory was available to fulfill the request.

37196 **EXAMPLES**

37197 None.

37198 **APPLICATION USAGE**37199 The *select()* and *poll()* functions can be used to determine when data is available to be received.37200 **RATIONALE**

37201 None.

37202 **FUTURE DIRECTIONS**

37203 None.

37204 **SEE ALSO**37205 *poll()*, *read()*, *recv()*, *recvmsg()*, *select()* (on page 1753)1 *send()*, *sendmsg()*, *sendto()*, *shutdown()*,  
37206 *socket()*, *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/socket.h>37207 **CHANGE HISTORY**

37208 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

## 37209 NAME

37210 recvmsg — receive a message from a socket

## 37211 SYNOPSIS

37212 #include &lt;sys/socket.h&gt;

37213 ssize\_t recvmsg(int *socket*, struct msghdr \**message*, int *flags*);

## 37214 DESCRIPTION

37215 The *recvmsg()* function receives a message from a connection-mode or connectionless-mode  
 37216 socket. It is normally used with connectionless-mode sockets because it permits the application  
 37217 to retrieve the source address of received data.

37218 The *recvmsg()* function takes the following arguments:

|       |                |                                                                                                                                                                                                                                                                                                                      |
|-------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 37219 | <i>socket</i>  | Specifies the socket file descriptor.                                                                                                                                                                                                                                                                                |
| 37220 | <i>message</i> | Points to a <b>msghdr</b> structure, containing both the buffer to store the source address and the buffers for the incoming message. The length and format of the address depend on the address family of the socket. The <i>msg_flags</i> member is ignored on input, but may contain meaningful values on output. |
| 37221 |                |                                                                                                                                                                                                                                                                                                                      |
| 37222 |                |                                                                                                                                                                                                                                                                                                                      |
| 37223 |                |                                                                                                                                                                                                                                                                                                                      |
| 37224 | <i>flags</i>   | Specifies the type of message reception. Values of this argument are formed by logically OR'ing zero or more of the following values:                                                                                                                                                                                |
| 37225 |                |                                                                                                                                                                                                                                                                                                                      |
| 37226 | MSG_OOB        | Requests out-of-band data. The significance and semantics of out-of-band data are protocol-specific.                                                                                                                                                                                                                 |
| 37227 |                |                                                                                                                                                                                                                                                                                                                      |
| 37228 | MSG_PEEK       | Peeks at the incoming message.                                                                                                                                                                                                                                                                                       |
| 37229 | MSG_WAITALL    | Requests that the function block until the full amount of data requested can be returned. The function may return a smaller amount of data if a signal is caught, if the connection is terminated, if MSG_PEEK was specified, or if an error is pending for the socket.                                              |
| 37230 |                |                                                                                                                                                                                                                                                                                                                      |
| 37231 |                |                                                                                                                                                                                                                                                                                                                      |
| 37232 |                |                                                                                                                                                                                                                                                                                                                      |
| 37233 |                |                                                                                                                                                                                                                                                                                                                      |

37234 The *recvmsg()* function receives messages from unconnected or connected sockets and shall  
 37235 return the length of the message.

37236 The *recvmsg()* function shall return the total length of the message. For message-based sockets,  
 37237 such as SOCK\_DGRAM and SOCK\_SEQPACKET, the entire message shall be read in a single  
 37238 operation. If a message is too long to fit in the supplied buffers, and MSG\_PEEK is not set in the  
 37239 *flags* argument, the excess bytes shall be discarded, and MSG\_TRUNC is set in the *msg\_flags*  
 37240 member of the **msghdr** structure. For stream-based sockets, such as SOCK\_STREAM, message  
 37241 boundaries shall be ignored. In this case, data is returned to the user as soon as it becomes  
 37242 available, and no data shall be discarded.

37243 If the MSG\_WAITALL flag is not set, data shall be returned only up to the end of the first  
 37244 message.

37245 If no messages are available at the socket and O\_NONBLOCK is not set on the socket's file  
 37246 descriptor, *recvmsg()* shall block until a message arrives. If no messages are available at the  
 37247 socket and O\_NONBLOCK is set on the socket's file descriptor, *recvmsg()* function shall fail and  
 37248 set *errno* to [EAGAIN] or [EWOULDBLOCK].

37249 In the **msghdr** structure, the *msg\_name* and *msg\_namelen* members specify the source address if  
 37250 the socket is unconnected. If the socket is connected, the *msg\_name* and *msg\_namelen* members  
 37251 are ignored. The *msg\_name* member may be a null pointer if no names are desired or required.  
 37252 The *msg\_iov* and *msg\_iovlen* fields are used to specify where the received data shall be stored.  
 37253 *msg\_iov* points to an array of **iovec** structures; *msg\_iovlen* shall be set to the dimension of this

37254 array. In each **iovec** structure, the *iov\_base* field specifies a storage area and the *iov\_len* field gives  
 37255 its size in bytes. Each storage area indicated by *msg\_iov* is filled with received data in turn until  
 37256 all of the received data is stored or all of the areas have been filled.

37257 Upon successful completion, the *msg\_flags* member of the message header is the bitwise-  
 37258 inclusive OR of all of the following flags that indicate conditions detected for the received  
 37259 message:

37260 MSG\_EOR End of record was received (if supported by the protocol).

37261 MSG\_OOB Out-of-band data was received.

37262 MSG\_TRUNC Normal data was truncated.

37263 MSG\_CTRUNC Control data was truncated.

#### 37264 RETURN VALUE

37265 Upon successful completion, *recvmsg()* shall return the length of the message in bytes. If no  
 37266 messages are available to be received and the peer has performed an orderly shutdown,  
 37267 *recvmsg()* shall return 0. Otherwise,  $-1$  shall be returned and *errno* set to indicate the error.

#### 37268 ERRORS

37269 The *recvmsg()* function shall fail if:

37270 [EAGAIN] or [EWOULDBLOCK]

37271 The socket's file descriptor is marked O\_NONBLOCK and no data is waiting  
 37272 to be received; or MSG\_OOB is set and no out-of-band data is available and  
 37273 either the socket's file descriptor is marked O\_NONBLOCK or the socket does  
 37274 not support blocking to await out-of-band data.

37275 [EBADF] The *socket* argument is not a valid open file descriptor.

37276 [ECONNRESET] A connection was forcibly closed by a peer.

37277 [EINTR] This function was interrupted by a signal before any data was available.

37278 [EINVAL] The sum of the *iov\_len* values overflows a **ssize\_t**, or the MSG\_OOB flag is set  
 37279 and no out-of-band data is available.

37280 [EMSGSIZE] The *msg\_iovlen* member of the **msg\_hdr** structure pointed to by *message* is less  
 37281 than or equal to 0, or is greater than {IOV\_MAX}.

37282 [ENOTCONN] A receive is attempted on a connection-mode socket that is not connected.

37283 [ENOTSOCK] The *socket* argument does not refer to a socket.

37284 [EOPNOTSUPP] The specified flags are not supported for this socket type.

37285 [ETIMEDOUT] The connection timed out during connection establishment, or due to a  
 37286 transmission timeout on active connection.

37287 The *recvmsg()* function may fail if:

37288 [EIO] An I/O error occurred while reading from or writing to the file system.

37289 [ENOBUFS] Insufficient resources were available in the system to perform the operation.

37290 [ENOMEM] Insufficient memory was available to fulfill the request.

37291 **EXAMPLES**

37292           None.

37293 **APPLICATION USAGE**37294           The *select()* and *poll()* functions can be used to determine when data is available to be received.37295 **RATIONALE**

37296           None.

37297 **FUTURE DIRECTIONS**

37298           None.

37299 **SEE ALSO**37300           *poll()*, *recv()*, *recvfrom()*, *select()*, *send()*, *sendmsg()*, *sendto()*, *shutdown()*, *socket()*, the Base |37301           Definitions volume of IEEE Std. 1003.1-200x, <**sys/socket.h**> |37302 **CHANGE HISTORY**

37303           First released in Issue 6. Derived from the XNS, Issue 5.2 specification. |

## 37304 NAME

37305 regcomp, regerror, regex, regfree — regular expression matching

## 37306 SYNOPSIS

37307 #include &lt;regex.h&gt;

```

37308 int regcomp(regex_t *restrict preg, const char *restrict pattern, int cflags);
37309 size_t regerror(int errcode, const regex_t *restrict preg,
37310 char *restrict errbuf, size_t errbuf_size);
37311 int regex(const regex_t *restrict preg, const char *restrict string,
37312 size_t nmatch, regmatch_t pmatch[restrict], int eflags);
37313 void regfree(regex_t *preg);

```

## 37314 DESCRIPTION

37315 These functions interpret *basic* and *extended* regular expressions as described in the Base  
 37316 Definitions volume of IEEE Std. 1003.1-200x, Chapter 9, Regular Expressions.

37317 The `regex_t` structure contains at least the following member:

37318

37319

37320

| Member Type | Member Name | Description                             |
|-------------|-------------|-----------------------------------------|
| size_t      | re_nsub     | Number of parenthesized subexpressions. |

37321 The `regmatch_t` structure contains at least the following members:

37322

37323

37324

37325

37326

| Member Type | Member Name | Description                                                                                |
|-------------|-------------|--------------------------------------------------------------------------------------------|
| regoff_t    | rm_so       | Byte offset from start of <i>string</i> to start of substring.                             |
| regoff_t    | rm_eo       | Byte offset from start of <i>string</i> of the first character after the end of substring. |

37327 The `regcomp()` function shall compile the regular expression contained in the string pointed to by  
 37328 the *pattern* argument and places the results in the structure pointed to by *preg*. The *cflags*  
 37329 argument is the bitwise-inclusive OR of zero or more of the following flags, which are defined in  
 37330 the header <regex.h>:

37331 REG\_EXTENDED     Use Extended Regular Expressions.

37332 REG\_ICASE        Ignore case in match. (See the Base Definitions volume of  
 37333 IEEE Std. 1003.1-200x, Chapter 9, Regular Expressions.)

37334 REG\_NOSUB        Report only success/fail in `regex()`.

37335 REG\_NEWLINE     Change the handling of <newline> characters, as described in the text.

37336 The default regular expression type for *pattern* is a Basic Regular Expression. The application can  
 37337 specify Extended Regular Expressions using the REG\_EXTENDED *cflags* flag.

37338 If the REG\_NOSUB flag was not set in *cflags*, then `regcomp()` shall set *re\_nsub* to the number of  
 37339 parenthesized subexpressions (delimited by "\(\)" in basic regular expressions or "( )" in  
 37340 extended regular expressions) found in *pattern*.

37341 The `regex()` function compares the null-terminated string specified by *string* with the compiled  
 37342 regular expression *preg* initialized by a previous call to `regcomp()`. If it finds a match, `regex()`  
 37343 shall return 0; otherwise, it shall return non-zero indicating either no match or an error. The  
 37344 *eflags* argument is the bitwise-inclusive OR of zero or more of the following flags, which are  
 37345 defined in the header <regex.h>:

37346 REG\_NOTBOL The first character of the string pointed to by *string* is not the beginning of the  
 37347 line. Therefore, the circumflex character ('^'), when taken as a special  
 37348 character, shall not match the beginning of *string*.

37349 REG\_NOTEOL The last character of the string pointed to by *string* is not the end of the line.  
 37350 Therefore, the dollar sign ('\$'), when taken as a special character, shall not  
 37351 match the end of *string*.

37352 If *nmatch* is 0 or REG\_NOSUB was set in the *flags* argument to *regcomp()*, then *regexec()* shall  
 37353 ignore the *pmatch* argument. Otherwise, the application shall ensure that the *pmatch* argument  
 37354 points to an array with at least *nmatch* elements, and *regexec()* shall fill in the elements of that  
 37355 array with offsets of the substrings of *string* that correspond to the parenthesized subexpressions  
 37356 of *pattern*: *pmatch[i].rm\_so* shall be the byte offset of the beginning and *pmatch[i].rm\_eo* shall be  
 37357 one greater than the byte offset of the end of substring *i*. (Subexpression *i* begins at the *i*th  
 37358 matched open parenthesis, counting from 1.) Offsets in *pmatch[0]* identify the substring that  
 37359 corresponds to the entire regular expression. Unused elements of *pmatch* up to *pmatch[nmatch-1]*  
 37360 shall be filled with -1. If there are more than *nmatch* subexpressions in *pattern* (*pattern* itself  
 37361 counts as a subexpression), then *regexec()* shall still do the match, but shall record only the first  
 37362 *nmatch* substrings.

37363 When matching a basic or extended regular expression, any given parenthesized subexpression  
 37364 of *pattern* might participate in the match of several different substrings of *string*, or it might not  
 37365 match any substring even though the pattern as a whole did match. The following rules are used  
 37366 to determine which substrings to report in *pmatch* when matching regular expressions:

- 37367 1. If subexpression *i* in a regular expression is not contained within another subexpression,  
 37368 and it participated in the match several times, then the byte offsets in *pmatch[i]* shall  
 37369 delimit the last such match.
- 37370 2. If subexpression *i* is not contained within another subexpression, and it did not participate  
 37371 in an otherwise successful match, the byte offsets in *pmatch[i]* shall be -1. A subexpression  
 37372 does not participate in the match when:

37373 ' \* ' or "\{\}" appears immediately after the subexpression in a basic regular  
 37374 expression, or ' \* ', ' ? ', or "{ }" appears immediately after the subexpression in an  
 37375 extended regular expression, and the subexpression did not match (matched 0 times)

37376 or:

37377 ' | ' is used in an extended regular expression to select this subexpression or another,  
 37378 and the other subexpression matched.

- 37379 3. If subexpression *i* is contained within another subexpression *j*, and *i* is not contained  
 37380 within any other subexpression that is contained within *j*, and a match of subexpression *j*  
 37381 is reported in *pmatch[j]*, then the match or non-match of subexpression *i* reported in  
 37382 *pmatch[i]* shall be as described in 1. and 2. above, but within the substring reported in  
 37383 *pmatch[j]* rather than the whole string. The offsets in *pmatch[i]* are still relative to the start  
 37384 of string.

37385 **Notes to Reviewers**37386 *This section with side shading will not appear in the final copy. - Ed.*37387 D1, XSH, ERN 283 proposes changing “but within” above to “but describing the substring  
37388 found within the substring reported in *pmatch[j]* rather than the whole string. The byte  
37389 offsets are relative to the whole string.”37390 4. If subexpression *i* is contained in subexpression *j*, and the byte offsets in *pmatch[j]* are  $-1$ ,  
37391 then the pointers in *pmatch[i]* shall also be  $-1$ .37392 5. If subexpression *i* matched a zero-length string, then both byte offsets in *pmatch[i]* shall be  
37393 the byte offset of the character or null terminator immediately following the zero-length  
37394 string.37395 If, when *regexec()* is called, the locale is different from when the regular expression was  
37396 compiled, the result is undefined.37397 If REG\_NEWLINE is not set in *cflags*, then a <newline> character in *pattern* or *string* shall be  
37398 treated as an ordinary character. If REG\_NEWLINE is set, then <newline> shall be treated as an  
37399 ordinary character except as follows:37400 1. A <newline> character in *string* shall not be matched by a period outside a bracket  
37401 expression or by any form of a non-matching list (see the Base Definitions volume of  
37402 IEEE Std. 1003.1-200x, Chapter 9, Regular Expressions).37403 2. A circumflex ('^') in *pattern*, when used to specify expression anchoring (see the Base  
37404 Definitions volume of IEEE Std. 1003.1-200x, Section 9.3.8, BRE Expression Anchoring),  
37405 shall match the zero-length string immediately after a <newline> in *string*, regardless of  
37406 the setting of REG\_NOTBOL.37407 3. A dollar sign ('\$') in *pattern*, when used to specify expression anchoring, shall match the  
37408 zero-length string immediately before a <newline> in *string*, regardless of the setting of  
37409 REG\_NOTEOL.37410 The *regfree()* function frees any memory allocated by *regcomp()* associated with *preg*.

37411 The following constants are defined as error return values:

37412 REG\_NOMATCH *regexec()* failed to match.

37413 REG\_BADPAT Invalid regular expression.

37414 REG\_ECOLLATE Invalid collating element referenced.

37415 REG\_ECTYPE Invalid character class type referenced.

37416 REG\_EESCAPE Trailing '\\' in pattern.

37417 REG\_ESUBREG Number in "\digit" invalid or in error.

37418 REG\_EBRACK "[ ]" imbalance.

37419 REG\_EPAREN "\(\)" or "()" imbalance.

37420 REG\_EBRACE "\{\}" imbalance.

37421 REG\_BADBR Content of "\{\}" invalid: not a number, number too large, more than  
37422 two numbers, first larger than second.

37423 REG\_ERANGE Invalid endpoint in range expression.

37424 REG\_ESPACE Out of memory.

37425 REG\_BADRPT '?', '\*', or '+' not preceded by valid regular expression.

37426 The *regerror()* function provides a mapping from error codes returned by *regcomp()* and  
 37427 *regexec()* to unspecified printable strings. It generates a string corresponding to the value of the  
 37428 *errcode* argument, which the application shall ensure is the last non-zero value returned by  
 37429 *regcomp()* or *regexec()* with the given value of *preg*. If *errcode* is not such a value, the content of  
 37430 the generated string is unspecified.

37431 If *preg* is a null pointer, but *errcode* is a value returned by a previous call to *regexec()* or *regcomp()*,  
 37432 the *regerror()* still generates an error string corresponding to the value of *errcode*, but it might not  
 37433 be as detailed under some implementations.

37434 If the *errbuf\_size* argument is not 0, *regerror()* shall place the generated string into the buffer of  
 37435 size *errbuf\_size* bytes pointed to by *errbuf*. If the string (including the terminating null) cannot fit  
 37436 in the buffer, *regerror()* shall truncate the string and null-terminates the result.

37437 If *errbuf\_size* is 0, *regerror()* shall ignore the *errbuf* argument, and return the size of the buffer  
 37438 needed to hold the generated string.

37439 If the *preg* argument to *regexec()* or *regfree()* is not a compiled regular expression returned by  
 37440 *regcomp()*, the result is undefined. A *preg* is no longer treated as a compiled regular expression  
 37441 after it is given to *regfree()*.

37442 **RETURN VALUE**

37443 Upon successful completion, the *regcomp()* function shall return 0. Otherwise, it shall return an  
 37444 integer value indicating an error as described in <regex.h>, and the content of *preg* is undefined.  
 37445 If a code is returned, the interpretation shall be as given in <regex.h>.

37446 If *regcomp()* detects an invalid RE, it may return REG\_BADPAT, or it may return one of the error  
 37447 codes that more precisely describes the error.

37448 Upon successful completion, the *regexec()* function shall return 0. Otherwise, it shall return  
 37449 REG\_NOMATCH to indicate no match.

37450 Upon successful completion, the *regerror()* function shall return the number of bytes needed to  
 37451 hold the entire generated string, including the null termination. If the return value is greater than  
 37452 *errbuf\_size*, the string returned in the buffer pointed to by *errbuf* has been truncated.

37453 The *regfree()* function shall return no value.

37454 **ERRORS**

37455 No errors are defined.

37456 **EXAMPLES**

```

37457 #include <regex.h>
37458 /*
37459 * Match string against the extended regular expression in
37460 * pattern, treating errors as no match.
37461 *
37462 * Return 1 for match, 0 for no match.
37463 */
37464 int
37465 match(const char *string, char *pattern)
37466 {
37467 int status;
37468 regex_t re;
```

```

37469 if (regcomp(&re, pattern, REG_EXTENDED|REG_NOSUB) != 0) {
37470 return(0); /* Report error. */
37471 }
37472 status = regexec(&re, string, (size_t) 0, NULL, 0);
37473 regfree(&re);
37474 if (status != 0) {
37475 return(0); /* Report error. */
37476 }
37477 return(1);
37478 }

```

37479 The following demonstrates how the REG\_NOTBOL flag could be used with *regexec()* to find all  
37480 substrings in a line that match a pattern supplied by a user. (For simplicity of the example, very  
37481 little error checking is done.)

```

37482 (void) regcomp (&re, pattern, 0);
37483 /* This call to regexec() finds the first match on the line. */
37484 error = regexec (&re, &buffer[0], 1, &pm, 0);
37485 while (error == 0) { /* While matches found. */
37486 /* Substring found between pm.rm_so and pm.rm_eo. */
37487 /* This call to regexec() finds the next match. */
37488 error = regexec (&re, buffer + pm.rm_eo, 1, &pm, REG_NOTBOL);
37489 }

```

#### 37490 APPLICATION USAGE

37491 An application could use:

```

37492 regerror(code, preg, (char *)NULL, (size_t)0)

```

37493 to find out how big a buffer is needed for the generated string, *malloc()* a buffer to hold the  
37494 string, and then call *regerror()* again to get the string. Alternatively, it could allocate a fixed,  
37495 static buffer that is big enough to hold most strings, and then use *malloc()* to allocate a larger  
37496 buffer if it finds that this is too small.

37497 To match a pattern as described in the Shell and Utilities volume of IEEE Std. 1003.1-200x,  
37498 Section 2.14, Pattern Matching Notation, use the *fnmatch()* function.

#### 37499 RATIONALE

37500 The *regmatch()* function must fill in all *nmatch* elements of *pmatch*, where *nmatch* and *pmatch* are  
37501 supplied by the application, even if some elements of *pmatch* do not correspond to  
37502 subexpressions in *pattern*. The application writer should note that there is probably no reason  
37503 for using a value of *nmatch* that is larger than *preg->re\_nsub+1*.

37504 The REG\_NEWLINE flag supports a use of RE matching that is needed in some applications like  
37505 text editors. In such applications, the user supplies an RE asking the application to find a line  
37506 that matches the given expression. An anchor in such an RE anchors at the beginning or end of  
37507 any line. Such an application can pass a sequence of <newline>-separated lines to *regexec()* as a  
37508 single long string and specify REG\_NEWLINE to *regcomp()* to get the desired behavior. The  
37509 application must ensure that there are no explicit <newline>s in *pattern* if it wants to ensure that  
37510 any match occurs entirely within a single line.

37511 The REG\_NEWLINE flag affects the behavior of *regexec()*, but it is in the *cflags* parameter to  
37512 *regcomp()* to allow flexibility of implementation. Some implementations will want to generate  
37513 the same compiled RE in *regcomp()* regardless of the setting of REG\_NEWLINE and have  
37514 *regexec()* handle anchors differently based on the setting of the flag. Other implementations will  
37515 generate different compiled REs based on the REG\_NEWLINE.

37516 The REG\_ICASE flag supports the operations taken by the *grep -i* option and the historical  
37517 implementations of *ex* and *vi*. Including this flag will make it easier for application code to be  
37518 written that does the same thing as these utilities.

37519 The substrings reported in *pmatch*[] are defined using offsets from the start of the string rather  
37520 than pointers. Since this is a new interface, there should be no impact on historical  
37521 implementations or applications, and offsets should be just as easy to use as pointers. The  
37522 change to offsets was made to facilitate future extensions in which the string to be searched is  
37523 presented to *regexexec*() in blocks, allowing a string to be searched that is not all in memory at  
37524 once.

37525 A new type **regoff\_t** is used for the elements of *pmatch*[] to ensure that the application can  
37526 represent either the largest possible array in memory (important for an application conforming  
37527 to the Shell and Utilities volume of IEEE Std. 1003.1-200x) or the largest possible file (important  
37528 for an application using the extension where a file is searched in chunks).

37529 The standard developers rejected the inclusion of a *regsub*() function that would be used to do  
37530 substitutions for a matched RE. While such a routine would be useful to some applications, its  
37531 utility would be much more limited than the matching function described here. Both RE parsing  
37532 and substitution are possible to implement without support other than that required by the  
37533 ISO C standard, but matching is much more complex than substituting. The only difficult part of  
37534 substitution, given the information supplied by *regexexec*(), is finding the next character in a string  
37535 when there can be multibyte characters. That is a much larger issue, and one that needs a more  
37536 general solution.

37537 The *errno* variable has not been used for error returns to avoid filling the *errno* name space for  
37538 this feature.

37539 The interface is defined so that the matched substrings *rm\_sp* and *rm\_ep* are in a separate  
37540 **regmatch\_t** structure instead of in **regex\_t**. This allows a single compiled RE to be used  
37541 simultaneously in several contexts; in *main*() and a signal handler, perhaps, or in multiple  
37542 threads of lightweight processes. (The *preg* argument to *regexexec*() is declared with type **const**, so  
37543 the implementation is not permitted to use the structure to store intermediate results.) It also  
37544 allows an application to request an arbitrary number of substrings from an RE. The number of  
37545 subexpressions in the RE is reported in *re\_nsub* in *preg*. With this change to *regexexec*(),  
37546 consideration was given to dropping the REG\_NOSUB flag since the user can now specify this  
37547 with a zero *nmatch* argument to *regexexec*(). However, keeping REG\_NOSUB allows an  
37548 implementation to use a different (perhaps more efficient) algorithm if it knows in *regcomp*()  
37549 that no subexpressions need be reported. The implementation is only required to fill in *pmatch* if  
37550 *nmatch* is not zero and if REG\_NOSUB is not specified. Note that the **size\_t** type, as defined in  
37551 the ISO C standard, is unsigned, so the description of *regexexec*() does not need to address  
37552 negative values of *nmatch*.

37553 REG\_NOTBOL was added to allow an application to do repeated searches for the same pattern  
37554 in a line. If the pattern contains a circumflex character that should match the beginning of a line,  
37555 then the pattern should only match when matched against the beginning of the line. Without  
37556 the REG\_NOTBOL flag, the application could rewrite the expression for subsequent matches,  
37557 but in the general case this would require parsing the expression. The need for REG\_NOTEOL is  
37558 not as clear; it was added for symmetry.

37559 The addition of the *regerror*() function addresses the historical need for portable application  
37560 programs to have access to error information more than “Function failed to compile/match your  
37561 RE for unknown reasons”.

37562 This interface provides for two different methods of dealing with error conditions. The specific  
37563 error codes (REG\_EBRACE, for example), defined in **<regex.h>**, allow an application to recover

37564 from an error if it is so able. Many applications, especially those that use patterns supplied by a  
 37565 user, will not try to deal with specific error cases, but will just use *regerror()* to obtain a human-  
 37566 readable error message to present to the user.

37567 The *regerror()* function uses a scheme similar to *confstr()* to deal with the problem of allocating  
 37568 memory to hold the generated string. The scheme used by *strerror()* in the ISO C standard was  
 37569 considered unacceptable since it creates difficulties for multi-threaded applications.

37570 The *preg* argument is provided to *regerror()* to allow an implementation to generate a more  
 37571 descriptive message than would be possible with *errcode* alone. An implementation might, for  
 37572 example, save the character offset of the offending character of the pattern in a field of *preg*, and  
 37573 then include that in the generated message string. The implementation may also ignore *preg*.

37574 A REG\_FILENAME flag was considered, but omitted. This flag caused *regexec()* to match  
 37575 patterns as described in the Shell and Utilities volume of IEEE Std. 1003.1-200x, Section 2.14,  
 37576 Pattern Matching Notation instead of REs. This service is now provided by the *fnmatch()*  
 37577 function.

37578 Notice that there is a difference in philosophy between the ISO POSIX-2:1993 standard and  
 37579 IEEE Std. 1003.1-200x in how to handle a bad regular expression. The ISO POSIX-2:1993  
 37580 standard says that many bad constructs produce undefined results, or that the interpretation is  
 37581 undefined. IEEE Std. 1003.1-200x, however, says that the interpretation of such REs is  
 37582 unspecified. The term “undefined” means that the action by the application is an error, of  
 37583 similar severity to passing a bad pointer to a function.

37584 The *regcomp()* and *regexec()* functions are required to accept any null-terminated string as the  
 37585 *pattern* argument. If the meaning of the string is undefined, the behavior of the function is  
 37586 unspecified. IEEE Std. 1003.1-200x does not specify how the functions will interpret the pattern;  
 37587 they might return error codes, or they might do pattern matching in some completely  
 37588 unexpected way, but they should not do something like abort the process.

#### 37589 FUTURE DIRECTIONS

37590 None.

#### 37591 SEE ALSO

37592 *fnmatch()*, *glob()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<regex.h>`,  
 37593 `<sys/types.h>`

#### 37594 CHANGE HISTORY

37595 First released in Issue 4. Derived from the ISO POSIX-2 standard.

#### 37596 Issue 5

37597 Moved from POSIX2 C-language Binding to BASE.

#### 37598 Issue 6

37599 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

37600 The following new requirements on POSIX implementations derive from alignment with the  
 37601 Single UNIX Specification:

- 37602 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
 37603 required for conforming implementations of previous POSIX specifications, it was not  
 37604 required for UNIX applications.

37605 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

37606 The REG\_ENOSYS constant is removed.

37607 The **restrict** keyword is added to the *regcomp()*, *regerror()*, and *regexec()* prototypes for  
 37608 alignment with the ISO/IEC 9899:1999 standard.



37609 **NAME**

37610 remainder, remainderf, remainderl — remainder function

37611 **SYNOPSIS**37612 XSI `#include <math.h>`37613 `double remainder(double x, double y);`37614 `float remainderf(float x, float y);`37615 `long double remainderl(long double x, long double y);`

37616

37617 **DESCRIPTION**

37618 These functions shall return the floating point remainder  $r=x-ny$  when  $y$  is non-zero. The value  
 37619  $n$  is the integral value nearest the exact value  $x/y$ . When  $|n-x/y|=1/2$ , the value  $n$  is chosen to be  
 37620 even.

37621 The behavior of *remainder()* is independent of the rounding mode.37622 **RETURN VALUE**37623 These functions shall return the floating point remainder  $r=x-ny$  when  $y$  is non-zero.37624 When  $y$  is 0, *remainder()* shall return NaN (or equivalent if available) and set *errno* to [EDOM].37625 If the value of  $x$  is  $\pm\text{Inf}$ , *remainder()* shall return NaN and set *errno* to [EDOM].37626 If  $x$  or  $y$  is NaN, then the function shall return NaN and *errno* may be set to [EDOM].37627 **ERRORS**

37628 These functions shall fail if:

37629 [EDOM] The  $y$  argument is 0 or the  $x$  argument is positive or negative infinity.

37630 These functions may fail if:

37631 [EDOM] The  $x$  or  $y$  argument is NaN.37632 **EXAMPLES**

37633 None.

37634 **APPLICATION USAGE**

37635 None.

37636 **RATIONALE**

37637 None.

37638 **FUTURE DIRECTIONS**

37639 None.

37640 **SEE ALSO**37641 *abs()*, *div()*, *ldiv()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<math.h>`37642 **CHANGE HISTORY**

37643 First released in Issue 4, Version 2.

37644 **Issue 5**

37645 Moved from X/OPEN UNIX extension to BASE.

37646 **Issue 6**37647 The *remainderf()* and *remainderl()* functions are added for alignment with the ISO/IEC 9899:1999  
 37648 standard.

37649 **NAME**

37650 remove — remove a file

37651 **SYNOPSIS**

37652 #include &lt;stdio.h&gt;

37653 int remove(const char \*path);

37654 **DESCRIPTION**

37655 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
37656 conflict between the requirements described here and the ISO C standard is unintentional. This  
37657 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

37658 The *remove()* function shall cause the file named by the path name pointed to by path to be no  
37659 longer accessible by that name. A subsequent attempt to open that file using that name shall fail,  
37660 unless it is created anew.

37661 CX If path does not name a directory, *remove(path)* shall be equivalent to *unlink(path)*.

37662 If path names a directory, *remove(path)* shall be equivalent to *rmdir(path)*.

37663 **RETURN VALUE**37664 CX Refer to *rmdir()* or *unlink()*.37665 **ERRORS**37666 CX Refer to *rmdir()* or *unlink()*.37667 **EXAMPLES**37668 **Removing Access to a File**37669 The following example shows how to remove access to a file named `/home/cnd/old_mods`.

```
37670 #include <stdio.h>
37671 int status;
37672 ...
37673 status = remove("/home/cnd/old_mods");
```

37674 **APPLICATION USAGE**

37675 None.

37676 **RATIONALE**

37677 None.

37678 **FUTURE DIRECTIONS**

37679 None.

37680 **SEE ALSO**37681 *rmdir()*, *unlink()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>37682 **CHANGE HISTORY**

37683 First released in Issue 3.

37684 Entry included for alignment with the POSIX.1-1988 standard and the ISO C standard.

37685 **Issue 4**

37686 All statements containing references to *unlink()* and *rmdir()* in the DESCRIPTION, RETURN  
37687 VALUE, and ERRORS sections are marked as extensions.

37688 The following changes are incorporated for alignment with the ISO C standard:

- 37689
  - The type of argument *path* is changed from **char\*** to **const char\***.
- 37690
  - The DESCRIPTION is expanded to describe the operation of *remove()* more completely.
- 37691 **Issue 6**
- 37692 Extensions beyond the ISO C standard are now marked.
- 37693 The following new requirements on POSIX implementations derive from alignment with the
- 37694 Single UNIX Specification:
  - The DESCRIPTION, RETURN VALUE, and ERRORS sections are updated so that if *path* is not a directory, *remove()* is equivalent to *unlink()*, and if it is a directory, it is equivalent to *rmdir()*.
- 37695
- 37696
- 37697

37698 **NAME**

37699       remque — remove an element from a queue

37700 **SYNOPSIS**

37701 xSI       #include <search.h>

37702       void remque(void \*element);

37703

37704 **DESCRIPTION**

37705       Refer to *insque()*.

37706 **NAME**

37707 remquo, remquof, remquol — remainder functions

37708 **SYNOPSIS**

37709 #include &lt;math.h&gt;

37710 double remquo(double x, double y, int \*quo);

37711 float remquof(float x, float y, int \*quo);

37712 long double remquol(long double x, long double y, int \*quo);

37713 **DESCRIPTION**

37714 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 37715 conflict between the requirements described here and the ISO C standard is unintentional. This  
 37716 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

37717 These functions shall compute the same remainder as the *remainder()*, *remainderf()*, and  
 37718 *remainderl()* functions, respectively. In the object pointed to by *quo* they store a value whose sign  
 37719 is the sign of  $x/y$  and whose magnitude is congruent modulo  $2^n$  to the magnitude of the integral  
 37720 quotient of  $x/y$ , where  $n$  is an implementation-defined integer greater than or equal to 3.

37721 An application wishing to check for error situations should set *errno* to 0 before calling these  
 37722 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

37723 **RETURN VALUE**37724 These functions shall return  $x \text{ REM } y$ .

37725 When  $y$  is 0, these functions shall return NaN (or equivalent if available) and set *errno* to  
 37726 [EDOM].

37727 If the value of  $x$  is  $\pm\text{Inf}$ , these functions shall return NaN and set *errno* to [EDOM].

37728 If  $x$  or  $y$  is NaN, then these functions shall return NaN and *errno* may be set to [EDOM].

37729 **ERRORS**

37730 These functions shall fail if:

37731 [EDOM] The  $y$  argument is 0 or the  $x$  argument is positive or negative infinity.

37732 These functions may fail if:

37733 [EDOM] The  $x$  or  $y$  argument is NaN.

37734 **EXAMPLES**

37735 None.

37736 **APPLICATION USAGE**

37737 None.

37738 **RATIONALE**

37739 These functions are intended for implementing argument reductions which can exploit a few  
 37740 low-order bits of the quotient. Note that  $x$  may be so large in magnitude relative to  $y$  that an  
 37741 exact representation of the quotient is not practical.

37742 **FUTURE DIRECTIONS**

37743 None.

37744 **SEE ALSO**37745 *remainder()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

37746 **CHANGE HISTORY**

37747 First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard.

37748 **NAME**

37749 rename — rename a file

37750 **SYNOPSIS**

37751 #include &lt;stdio.h&gt;

37752 int rename(const char \*old, const char \*new);

37753 **DESCRIPTION**

37754 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 37755 conflict between the requirements described here and the ISO C standard is unintentional. This  
 37756 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

37757 The *rename()* function shall change the name of a file. The *old* argument points to the path name  
 37758 of the file to be renamed. The *new* argument points to the new path name of the file.

37759 cx If the *old* argument and the *new* argument both refer to, and both link to, the same existing file,  
 37760 *rename()* shall return successfully and perform no other action.

37761 If the *old* argument points to the path name of a file that is not a directory, the application shall  
 37762 ensure that the *new* argument does not point to the path name of a directory. If the link named  
 37763 by the *new* argument exists, it shall be removed and *old* renamed to *new*. In this case, a link  
 37764 named *new* shall remain visible to other processes throughout the renaming operation and refer  
 37765 either to the file referred to by *new* or *old* before the operation began. Write access permission is  
 37766 required for both the directory containing *old* and the directory containing *new*.

37767 If the *old* argument points to the path name of a directory, the application shall ensure that the  
 37768 *new* argument does not point to the path name of a file that is not a directory. If the directory  
 37769 named by the *new* argument exists, it shall be removed and *old* renamed to *new*. In this case, a  
 37770 link named *new* shall exist throughout the renaming operation and shall refer either to the  
 37771 directory referred to by *new* or *old* before the operation began. If *new* names an existing directory,  
 37772 the application shall ensure that it is an empty directory.

37773 If either the *old* or the *new* arguments name a symbolic link, *rename()* shall operate on the  
 37774 symbolic link itself, and shall not resolve the last component of the argument. If *old* points to a  
 37775 path name that names a symbolic link, the symbolic link shall be renamed. If *new* points to a  
 37776 path name that names a symbolic link, the symbolic link shall be removed.

37777 The application shall ensure that the *new* path name does not contain a path prefix that names  
 37778 *old*. Write access permission is required for the directory containing *old* and the directory  
 37779 containing *new*. If the *old* argument points to the path name of a directory, write access  
 37780 permission may be required for the directory named by *old*, and, if it exists, the directory named  
 37781 by *new*.

37782 If the link named by the *new* argument exists and the file's link count becomes 0 when it is  
 37783 removed and no process has the file open, the space occupied by the file shall be freed and the  
 37784 file shall no longer be accessible. If one or more processes have the file open when the last link is  
 37785 removed, the link shall be removed before *rename()* returns, but the removal of the file contents  
 37786 shall be postponed until all references to the file are closed.

37787 Upon successful completion, *rename()* shall mark for update the *st\_ctime* and *st\_mtime* fields of  
 37788 the parent directory of each file.

37789 If the *rename()* function fails for any reason other than [EIO], any file named by *new* shall be  
 37790 unaffected.

## 37791 RETURN VALUE

37792 CX Upon successful completion, *rename()* shall return 0; otherwise, -1 shall be returned, *errno* shall  
 37793 be set to indicate the error, and neither the file named by *old* nor the file named by *new* shall be  
 37794 changed or created.

## 37795 ERRORS

37796 The *rename()* function shall fail if:

37797 CX [EACCES] A component of either path prefix denies search permission; or one of the  
 37798 directories containing *old* or *new* denies write permissions; or, write  
 37799 permission is required and is denied for a directory pointed to by the *old* or  
 37800 *new* arguments.

37801 CX [EBUSY] The directory named by *old* or *new* is currently in use by the system or another  
 37802 process, and the implementation considers this an error.

37803 CX [EEXIST] or [ENOTEMPTY]  
 37804 The link named by *new* is a directory that is not an empty directory.

37805 CX [EINVAL] The *new* directory path name contains a path prefix that names the *old*  
 37806 directory.

37807 CX [EIO] A physical I/O error has occurred.

37808 CX [EISDIR] The *new* argument points to a directory and the *old* argument points to a file  
 37809 that is not a directory.

37810 CX [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 37811 argument.

37812 CX [EMLINK] The file named by *old* is a directory, and the link count of the parent directory  
 37813 of *new* would exceed {LINK\_MAX}.

37814 CX [ENAMETOOLONG]  
 37815 The length of the *old* or *new* argument exceeds {PATH\_MAX} or a path name  
 37816 component is longer than {NAME\_MAX}.

37817 CX [ENOENT] The link named by *old* does not name an existing file, or either *old* or *new*  
 37818 points to an empty string.

37819 CX [ENOSPC] The directory that would contain *new* cannot be extended.

37820 CX [ENOTDIR] A component of either path prefix is not a directory; or the *old* argument  
 37821 names a directory and *new* argument names a non-directory file.

37822 XSI [EPERM] or [EACCES]  
 37823 The S\_ISVTX flag is set on the directory containing the file referred to by *old*  
 37824 and the caller is not the file owner, nor is the caller the directory owner, nor  
 37825 does the caller have appropriate privileges; or *new* refers to an existing file, the  
 37826 S\_ISVTX flag is set on the directory containing this file, and the caller is not  
 37827 the file owner, nor is the caller the directory owner, nor does the caller have  
 37828 appropriate privileges.

37829 CX [EROFS] The requested operation requires writing in a directory on a read-only file  
 37830 system.

37831 CX [EXDEV] The links named by *new* and *old* are on different file systems and the  
 37832 implementation does not support links between file systems.

37833 The *rename()* function may fail if:

|       |     |                |                                                                                                                                                            |
|-------|-----|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 37834 | XSI | [EBUSY]        | The file named by the <i>old</i> or <i>new</i> arguments is a named STREAM.                                                                                |
| 37835 |     | [ELOOP]        | More than {SYMLOOP_MAX} symbolic links were encountered during resolution of the <i>path</i> argument.                                                     |
| 37836 |     |                |                                                                                                                                                            |
| 37837 | CX  | [ENAMETOOLONG] |                                                                                                                                                            |
| 37838 |     |                | As a result of encountering a symbolic link in resolution of the <i>path</i> argument, the length of the substituted path name string exceeded {PATH_MAX}. |
| 37839 |     |                |                                                                                                                                                            |
| 37840 | CX  | [ETXTBSY]      | The file to be renamed is a pure procedure (shared text) file that is being executed.                                                                      |
| 37841 |     |                |                                                                                                                                                            |

37842 **EXAMPLES**37843 **Renaming a File**

37844 The following example shows how to rename a file named `/home/cnd/mod1` to  
37845 `/home/cnd/mod2`.

```
37846 #include <stdio.h>
37847 int status;
37848 ...
37849 status = rename("/home/cnd/mod1", "/home/cnd/mod2");
```

37850 **APPLICATION USAGE**

37851 None.

37852 **RATIONALE**

37853 This `rename()` function is equivalent for regular files to that defined by the ISO C standard. Its  
37854 inclusion here expands that definition to include actions on directories and specifies behavior  
37855 when the *new* parameter names a file that already exists. That specification requires that the  
37856 action of the function be atomic.

37857 One of the reasons for introducing this function was to have a means of renaming directories  
37858 while permitting implementations to prohibit the use of `link()` and `unlink()` with directories,  
37859 thus constraining links to directories to those made by `mkdir()`.

37860 The specification that if *old* and *new* refer to the same file is intended to guarantee that:

```
37861 rename("x", "x");
37862 does not remove the file.
```

37863 Renaming dot or dot-dot is prohibited in order to prevent cyclical file system paths.

37864 See also the descriptions of [ENOTEMPTY] and [ENAMETOOLONG] in `rmdir()` and [EBUSY] in  
37865 `unlink()`. For a discussion of [EXDEV], see `link()`.

37866 **FUTURE DIRECTIONS**

37867 None.

37868 **SEE ALSO**

37869 `link()`, `rmdir()`, `symlink()`, `unlink()`, the Base Definitions volume of IEEE Std. 1003.1-200x,  
37870 `<stdio.h>`

37871 **CHANGE HISTORY**

37872 First released in Issue 3.

37873 Entry included for alignment with the POSIX.1-1988 standard.

37874 **Issue 4**

37875 The [EMLINK] error is added to the ERRORS section.

37876 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 37877
- The type of arguments *old* and *new* are changed from **char\*** to **const char\***.
  - The RETURN VALUE section now states that if an error occurs, neither file is changed or created.
- 37878  
37879

37880 The following change is incorporated for alignment with the FIPS requirements:

- 37881
- In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path name component is larger than {NAME\_MAX}, is now defined as mandatory and marked as an extension.
- 37882  
37883

37884 **Issue 4, Version 2**

37885 The following changes are made for X/OPEN UNIX conformance:

- 37886
- The DESCRIPTION is updated to indicate the results of naming a symbolic link in either *old* or *new*.
  - In the ERRORS section, [EIO] is added to indicate that a physical I/O error has occurred, [ELOOP] to indicate that too many symbolic links were encountered during path name resolution, and [EPERM] or [EACCES] to indicate a permission check failure when operating on directories with S\_ISVTX set.
  - In the ERRORS section, a second [ENAMETOOLONG] condition is defined that may report excessive length of an intermediate result of path name resolution of a symbolic link.
- 37887  
37888  
37889  
37890  
37891  
37892  
37893

37894 **Issue 5**

37895 The [EBUSY] error is added to the “may fail” part of the ERRORS section.

37896 **Issue 6**

37897 Extensions beyond the ISO C standard are now marked.

37898 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 37899
- The [ENAMETOOLONG] error is restored as an error dependent on \_POSIX\_NO\_TRUNC. This is since behavior may vary from one file system to another.
- 37900

37901 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 37902
- The [EIO] mandatory error condition is added.
  - The [ELOOP] mandatory error condition is added.
  - A second [ENAMETOOLONG] is added as an optional error condition.
  - The [ETXTBSY] optional error condition is added.
- 37903  
37904  
37905  
37906

37907 The following changes were made to align with the IEEE P1003.1a draft standard:

- 37908
- Details are added regarding the treatment of symbolic links.
  - The [ELOOP] optional error condition is added.
- 37909

37910 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

37911 **NAME**

37912           rewind — reset file position indicator in a stream

37913 **SYNOPSIS**

37914           #include <stdio.h>

37915           void rewind(FILE \**stream*);

37916 **DESCRIPTION**

37917 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
37918           conflict between the requirements described here and the ISO C standard is unintentional. This  
37919           volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

37920           The call:

37921           rewind(*stream*)

37922           shall be equivalent to:

37923           (void) fseek(*stream*, 0L, SEEK\_SET)

37924           except that *rewind()* also clears the error indicator.

37925           Because *rewind()* does not return a value, an application wishing to detect errors should clear  
37926           *errno*, then call *rewind()*, and if *errno* is non-zero, assume an error has occurred.

37927 **RETURN VALUE**

37928           The *rewind()* function shall return no value.

37929 **ERRORS**

37930 **CX**       Refer to *fseek()* with the exception of [EINVAL] which does not apply.

37931 **EXAMPLES**

37932           None.

37933 **APPLICATION USAGE**

37934           None.

37935 **RATIONALE**

37936           None.

37937 **FUTURE DIRECTIONS**

37938           None.

37939 **SEE ALSO**

37940           *fseek()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>

37941 **CHANGE HISTORY**

37942           First released in Issue 1. Derived from Issue 1 of the SVID.

37943 **Issue 6**

37944           Extensions beyond the ISO C standard are now marked.

37945 **NAME**

37946       rewinddir — reset position of directory stream to the beginning of a directory

37947 **SYNOPSIS**

37948       #include <dirent.h>

37949       void rewinddir(DIR \*dirp);

37950 **DESCRIPTION**

37951       The *rewinddir()* function resets the position of the directory stream to which *dirp* refers to the  
37952       beginning of the directory. It shall also cause the directory stream to refer to the current state of  
37953       the corresponding directory, as a call to *opendir()* would have done. If *dirp* does not refer to a  
37954       directory stream, the effect is undefined.

37955       After a call to the *fork()* function, either the parent or child (but not both) may continue  
37956 XSI       processing the directory stream using *readdir()*, *rewinddir()*, or *seekdir()*. If both the parent and  
37957       child processes use these functions, the result is undefined.

37958 **RETURN VALUE**

37959       The *rewinddir()* function shall not return a value.

37960 **ERRORS**

37961       No errors are defined.

37962 **EXAMPLES**

37963       None.

37964 **APPLICATION USAGE**

37965       The *rewinddir()* function should be used in conjunction with *opendir()*, *readdir()*, and *closedir()* to  
37966       examine the contents of the directory. This method is recommended for portability.

37967 **RATIONALE**

37968       None.

37969 **FUTURE DIRECTIONS**

37970       None.

37971 **SEE ALSO**

37972       *closedir()*, *opendir()*, *readdir()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <dirent.h>  
37973       <sys/types.h>

37974 **CHANGE HISTORY**

37975       First released in Issue 2.

37976 **Issue 4**

37977       The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
37978       XSI-conformant systems.

37979       The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 37980       • The last paragraph of the DESCRIPTION, describing a restriction after a *fork()* function, is  
37981       added.

37982 **Issue 6**

37983       In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

37984       The following new requirements on POSIX implementations derive from alignment with the  
37985       Single UNIX Specification:

- 37986       • The requirement to include <sys/types.h> has been removed. Although <sys/types.h> was  
37987       required for conforming implementations of previous POSIX specifications, it was not

37988

required for UNIX applications.

37989 **NAME**37990 rindex — character string operations (**LEGACY**)37991 **SYNOPSIS**37992 XSI `#include <strings.h>`37993 `char *rindex(const char *s, int c);`

37994

37995 **DESCRIPTION**37996 The *rindex()* function is identical to *strchr()*.37997 **RETURN VALUE**37998 Refer to *strchr()*.37999 **ERRORS**38000 Refer to *strchr()*.38001 **EXAMPLES**

38002 None.

38003 **APPLICATION USAGE**38004 *strchr()* is preferred over this function.38005 For maximum portability, it is recommended to replace the function call to *rindex()* as follows:38006 `#define rindex(a,b) strchr((a),(b))`38007 **RATIONALE**

38008 None.

38009 **FUTURE DIRECTIONS**

38010 This function may be withdrawn in a future version.

38011 **SEE ALSO**38012 *strchr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<strings.h>`38013 **CHANGE HISTORY**

38014 First released in Issue 4, Version 2.

38015 **Issue 5**

38016 Moved from X/OPEN UNIX extension to BASE.

38017 **Issue 6**

38018 This function is marked LEGACY.

38019 **NAME**

38020 rint, rintf, rintl — round-to-nearest integral value

38021 **SYNOPSIS**

38022 #include <math.h>

38023 double rint(double x);

38024 float rintf(float x);

38025 long double rintl(long double x);

38026 **DESCRIPTION**

38027 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
38028 conflict between the requirements described here and the ISO C standard is unintentional. This  
38029 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

38030 These functions shall return the integral value (represented as a **double**) nearest  $x$  in the  
38031 direction of the current rounding mode. The current rounding mode is implementation-defined.

38032 If the current rounding mode rounds toward negative infinity, then *rint()* is identical to *floor()*.  
38033 If the current rounding mode rounds toward positive infinity, then *rint()* is identical to *ceil()*.

38034 These functions differ from the *nearbyint()*, *nearbyintf()*, and *nearbyintl()* functions only in that  
38035 they may raise the inexact floating-point exception if the result differs in value from the  
38036 argument.

38037 **RETURN VALUE**

38038 Upon successful completion, these functions shall return the integer (represented as a double  
38039 precision number) nearest  $x$  in the direction of the current rounding mode.

38040 When  $x$  is  $\pm\text{Inf}$ , *rint()* shall return  $x$ .

38041 If the value of  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].

38042 **ERRORS**

38043 These functions may fail if:

38044 [EDOM] The  $x$  argument is NaN.

38045 **EXAMPLES**

38046 None.

38047 **APPLICATION USAGE**

38048 None.

38049 **RATIONALE**

38050 None.

38051 **FUTURE DIRECTIONS**

38052 None.

38053 **SEE ALSO**

38054 *abs()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

38055 **CHANGE HISTORY**

38056 First released in Issue 4, Version 2.

38057 **Issue 5**

38058 Moved from X/OPEN UNIX extension to BASE.

38059 **Issue 6**

38060 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 38061
- The *rintf()* and *rintl()* functions are added.
- 38062
- The *rint()* function is no longer marked XSI as it is part of the ISO/IEC 9899:1999 standard.

38063 **NAME**

38064           rmdir — remove a directory

38065 **SYNOPSIS**

38066           #include &lt;unistd.h&gt;

38067           int rmdir(const char \*path);

38068 **DESCRIPTION**38069           The *rmdir()* function shall remove a directory whose name is given by *path*. The directory is  
38070 removed only if it is an empty directory.38071           If the directory is the root directory or the current working directory of any process, it is  
38072 unspecified whether the function succeeds, or whether it shall fail and set *errno* to [EBUSY].38073           If *path* names a symbolic link, then *rmdir()* shall fail and set *errno* to [ENOTDIR].38074           If the *path* argument refers to a path whose final component is either dot or dot-dot, *rmdir()* shall  
38075 fail.38076           If the directory's link count becomes 0 and no process has the directory open, the space occupied  
38077 by the directory shall be freed and the directory shall no longer be accessible. If one or more  
38078 processes have the directory open when the last link is removed, the dot and dot-dot entries, if  
38079 present, are removed before *rmdir()* returns and no new entries may be created in the directory,  
38080 but the directory is not removed until all references to the directory are closed.38081           If the directory is not an empty directory, *rmdir()* shall fail and set *errno* to [EEXIST] or  
38082 [ENOTEMPTY].38083           Upon successful completion, the *rmdir()* function shall mark for update the *st\_ctime* and  
38084 *st\_mtime* fields of the parent directory.38085 **RETURN VALUE**38086           Upon successful completion, the function *rmdir()* shall return 0. Otherwise, -1 shall be returned,  
38087 and *errno* set to indicate the error. If -1 is returned, the named directory shall not be changed.38088 **ERRORS**38089           The *rmdir()* function shall fail if:38090           [EACCES]           Search permission is denied on a component of the path prefix, or write  
38091 permission is denied on the parent directory of the directory to be removed.38092           [EBUSY]           The directory to be removed is currently in use by the system or some process  
38093 and the implementation considers this to be an error.

38094           [EEXIST] or [ENOTEMPTY]

38095                               The *path* argument names a directory that is not an empty directory, or there  
38096 are hard links to the directory other than dot or a single entry in dot-dot.38097           [EINVAL]           The *path* argument contains a last component that is dot.

38098           [EIO]              A physical I/O error has occurred.

38099           [ELOOP]           A loop exists in symbolic links encountered during resolution of the *path*  
38100 argument.

38101           [ENAMETOOLONG]

38102                               The length of the *path* argument exceeds {PATH\_MAX} or a path name  
38103 component is longer than

38104           NAME\_MAX

|           |                     |                                                                                        |
|-----------|---------------------|----------------------------------------------------------------------------------------|
| 38105     | [ENOENT]            | A component of <i>path</i> does not name an existing file, or the <i>path</i> argument |
| 38106     |                     | names a nonexistent directory or points to an empty string.                            |
| 38107     | [ENOTDIR]           | A component of <i>path</i> is not a directory.                                         |
| 38108 XSI | [EPERM] or [EACCES] |                                                                                        |
| 38109     |                     | The S_ISVTX flag is set on the parent directory of the directory to be removed         |
| 38110     |                     | and the caller is not the owner of the directory to be removed, nor is the caller      |
| 38111     |                     | the owner of the parent directory, nor does the caller have the appropriate            |
| 38112     |                     | privileges.                                                                            |
| 38113     | [EROFS]             | The directory entry to be removed resides on a read-only file system.                  |
| 38114     |                     | The <i>rmdir()</i> function may fail if:                                               |
| 38115     | [ELOOP]             | More than {SYMLOOP_MAX} symbolic links were encountered during                         |
| 38116     |                     | resolution of the <i>path</i> argument.                                                |
| 38117     | [ENAMETOOLONG]      |                                                                                        |
| 38118     |                     | As a result of encountering a symbolic link in resolution of the <i>path</i> argument, |
| 38119     |                     | the length of the substituted path name string exceeded {PATH_MAX}.                    |

### 38120 EXAMPLES

#### 38121 Removing a Directory

38122 The following example shows how to remove a directory named `/home/cnd/mod1`.

```
38123 #include <unistd.h>
38124 int status;
38125 ...
38126 status = rmdir("/home/cnd/mod1");
```

### 38127 APPLICATION USAGE

38128 None.

### 38129 RATIONALE

38130 The *rmdir()* and *rename()* functions originated in 4.2 BSD, and they used [ENOTEMPTY] for the  
 38131 condition when the directory to be removed does not exist or *new* already exists. When the 1984  
 38132 /usr/group standard was published, it contained [EEXIST] instead. When these functions were  
 38133 adopted into System V, the 1984 /usr/group standard was used as a reference. Therefore, several  
 38134 existing applications and implementations support/use both forms, and no agreement could be  
 38135 reached on either value. All implementations are required to supply both [EEXIST] and  
 38136 [ENOTEMPTY] in `<errno.h>` with distinct values, so that applications can use both values in C-  
 38137 language **case** statements.

38138 The meaning of deleting *pathname/dot* is unclear, because the name of the file (directory) in the  
 38139 parent directory to be removed is not clear, particularly in the presence of multiple links to a  
 38140 directory.

38141 IEEE Std. 1003.1-200x was silent with regard to the behavior of *rmdir()* when there are multiple  
 38142 hard links to the directory being removed. The requirement to set *errno* to [EEXIST] or  
 38143 [ENOTEMPTY] clarifies the behavior in this case.

38144 If the process's home directory is being removed, that should be an allowed error.

38145 Virtually all existing implementations detect [ENOTEMPTY] or the case of dot-dot. The text in  
 38146 Section 2.3 (on page 521) about returning any one of the possible errors permits that behavior to  
 38147 continue. The [ELOOP] error may be returned if more than {SYMLOOP\_MAX} symbolic links

38148 are encountered during resolution of the *path* argument.

#### 38149 FUTURE DIRECTIONS

38150 None.

#### 38151 SEE ALSO

38152 *mkdir()*, *remove()*, *unlink()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

#### 38153 CHANGE HISTORY

38154 First released in Issue 3.

38155 Entry included for alignment with the POSIX.1-1988 standard.

#### 38156 Issue 4

38157 The <**unistd.h**> header is added to the SYNOPSIS section.

38158 The [ENAMETOOLONG] description is amended.

38159 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 38160 • The type of argument *path* is changed from **char\*** to **const char\***.
- 38161 • The DESCRIPTION is expanded to indicate that, if the directory is a root directory or a
- 38162 current working directory, it is unspecified whether the function succeeds, or whether it fails
- 38163 and sets *errno* to [EBUSY]. In Issue 3, the behavior under these circumstances was defined as
- 38164 implementation-defined.
- 38165 • The RETURN VALUE section is expanded to direct that if  $-1$  is returned, the directory is not
- 38166 changed.

38167 The following change is incorporated for alignment with the FIPS requirements:

- 38168 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path
- 38169 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as
- 38170 an extension.

#### 38171 Issue 4, Version 2

38172 The following changes are made for X/OPEN UNIX conformance:

- 38173 • The DESCRIPTION is updated to indicate the results of naming a symbolic link in *path*.
- 38174 • In the ERRORS section, [EIO] is added to indicate that a physical I/O error has occurred,
- 38175 [ELOOP] to indicate that too many symbolic links were encountered during path name
- 38176 resolution, and [EPERM] or [EACCES] to indicate a permission check failure when operating
- 38177 on directories with S\_ISVTX set.
- 38178 • In the ERRORS section, a second [ENAMETOOLONG] condition is defined that may report
- 38179 excessive length of an intermediate result of path name resolution of a symbolic link.

#### 38180 Issue 6

38181 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 38182 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.
- 38183 This is since behavior may vary from one file system to another.

38184 The following new requirements on POSIX implementations derive from alignment with the

38185 Single UNIX Specification:

- 38186 • The DESCRIPTION is updated to indicate the results of naming a symbolic link in *path*.
- 38187 • The [EIO] mandatory error condition is added.

- 38188           • The [ELOOP] mandatory error condition is added.
- 38189           • A second [ENAMETOOLONG] is added as an optional error condition.
- 38190       The following changes were made to align with the IEEE P1003.1a draft standard:
- 38191           • The [ELOOP] optional error condition is added.

38192 **NAME**

38193 round, roundf, roundl — round to nearest integer value in floating-point format

38194 **SYNOPSIS**

38195 #include <math.h>

38196 double round(double x);

38197 float roundf(float x);

38198 long double roundl(long double x);

38199 **DESCRIPTION**

38200 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
38201 conflict between the requirements described here and the ISO C standard is unintentional. This  
38202 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

38203 These functions shall round their argument to the nearest integer value in floating-point format,  
38204 rounding halfway cases away from zero, regardless of the current rounding direction.

38205 An application wishing to check for error situations should set *errno* to 0 before calling these  
38206 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

38207 **RETURN VALUE**

38208 Upon successful completion, these functions shall return the rounded integer value.

38209 If *x* is  $\pm\text{Inf}$ , these functions shall return *x*.

38210 If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

38211 **ERRORS**

38212 These functions may fail if:

38213 [EDOM] The value of *x* is NaN.

38214 **EXAMPLES**

38215 None.

38216 **APPLICATION USAGE**

38217 None.

38218 **RATIONALE**

38219 None.

38220 **FUTURE DIRECTIONS**

38221 None.

38222 **SEE ALSO**

38223 The Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

38224 **CHANGE HISTORY**

38225 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

38226 **NAME**

38227 scalb — load exponent of a radix-independent floating-point number

38228 **SYNOPSIS**38229 XSI `#include <math.h>`38230 `double scalb(double x, double n);`

38231

38232 **DESCRIPTION**

38233 The *scalb()* function shall compute  $x \cdot r^n$ , where  $r$  is the radix of the machine's floating point  
 38234 arithmetic. When  $r$  is 2, *scalb()* is equivalent to *ldexp()*. The value of  $r$  is FLT\_RADIX which is  
 38235 defined in `<float.h>`.

38236 An application wishing to check for error situations should set *errno* to 0 before calling *scalb()*. If  
 38237 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

38238 **RETURN VALUE**38239 Upon successful completion, the *scalb()* function shall return  $x \cdot r^n$ .

38240 If the correct value would overflow, *scalb()* shall return  $\pm$ HUGE\_VAL (according to the sign of  $x$ )  
 38241 and set *errno* to [ERANGE].

38242 If the correct value would underflow, *scalb()* shall return 0 and set *errno* to [ERANGE].

38243 The *scalb()* function shall return  $x$  when  $x$  is  $\pm$ Inf.

38244 If  $x$  or  $n$  is NaN, then *scalb()* shall return NaN and may set *errno* to [EDOM].

38245 **ERRORS**38246 The *scalb()* function shall fail if:

38247 [ERANGE] The correct value would overflow or underflow.

38248 The *scalb()* function may fail if:

38249 [EDOM] The  $x$  or  $n$  argument is NaN.

38250 **EXAMPLES**

38251 None.

38252 **APPLICATION USAGE**

38253 None.

38254 **RATIONALE**

38255 None.

38256 **FUTURE DIRECTIONS**

38257 None.

38258 **SEE ALSO**

38259 *ilogb()*, *ldexp()*, *logb()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<float.h>`,  
 38260 `<math.h>`

38261 **CHANGE HISTORY**

38262 First released in Issue 4, Version 2.

38263 **Issue 5**

38264 Moved from X/OPEN UNIX extension to BASE.

38265 The DESCRIPTION is updated to indicate how an application should check for an error. This  
 38266 text was previously published in the APPLICATION USAGE section.

38267 **NAME**

38268 `scalbln, scalblnf, scalblnl scalbn, scalbnf, scalbnl`, — compute exponent using `FLT_RADIX`

38269 **SYNOPSIS**

38270 `#include <math.h>`

38271 `double scalbln(double x, long n);`  
 38272 `float scalblnf(float x, long n);`  
 38273 `long double scalblnl(long double x, long n);`  
 38274 `double scalbn(double x, int n);`  
 38275 `float scalbnf(float x, int n);`  
 38276 `long double scalbnl(long double x, int n);`

38277 **DESCRIPTION**

38278 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 38279 conflict between the requirements described here and the ISO C standard is unintentional. This  
 38280 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

38281 These functions shall compute  $x * FLT\_RADIX^n$  efficiently, not normally by computing  
 38282  $FLT\_RADIX^n$  explicitly.

38283 An application wishing to check for error situations should set *errno* to 0 before calling these  
 38284 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

38285 **RETURN VALUE**

38286 Upon successful completion, these functions shall return  $x * FLT\_RADIX^n$ .

38287 If the correct value would overflow, these functions shall return  $\pm HUGE\_VAL$  (according to the  
 38288 sign of *x*) and set *errno* to [ERANGE].

38289 If the correct value would underflow, these functions shall return 0 and set *errno* to [ERANGE].

38290 These functions shall return *x* when *x* is  $\pm Inf$ .

38291 If *x* or *n* is NaN, then these functions shall return NaN and may set *errno* to [EDOM].

38292 **ERRORS**

38293 These functions shall fail if:

38294 [ERANGE] The correct value would overflow or underflow.

38295 These functions may fail if:

38296 [EDOM] The *x* or *n* argument is NaN.

38297 **EXAMPLES**

38298 None.

38299 **APPLICATION USAGE**

38300 None.

38301 **RATIONALE**

38302 These functions are named so as to avoid conflicting with the Single UNIX Specification, which  
 38303 has a *scalb()* function whose second argument is **double** instead of **int**. The *scalb()* function is  
 38304 not part of ISO C standard. These functions, whose second parameter has type **long**, is provided  
 38305 because the factor required to scale from the smallest positive floating-point value to the largest  
 38306 finite one, on many implementations, is too large to represent in the minimum-width **int** format.

38307 **FUTURE DIRECTIONS**

38308           None.

38309 **SEE ALSO**

38310           *scalb()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

38311 **CHANGE HISTORY**

38312           First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard.

38313 **NAME**

38314           scanf — convert formatted input

38315 **SYNOPSIS**

38316           #include &lt;stdio.h&gt;

38317           int scanf(const char \*restrict *format*, ... );38318 **DESCRIPTION**38319           Refer to *fscanf()*.

38320 **NAME**

38321 sched\_get\_priority\_max, sched\_get\_priority\_min — get priority limits (**REALTIME**)

38322 **SYNOPSIS**

```
38323 PS #include <sched.h>
```

```
38324 int sched_get_priority_max(int policy);
```

```
38325 int sched_get_priority_min(int policy);
```

```
38326
```

38327 **DESCRIPTION**

38328 The *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()* functions return the appropriate  
38329 maximum or minimum, respectively, for the scheduling policy specified by *policy*.

38330 The value of *policy* is one of the scheduling policy values defined in <**sched.h**>.

38331 **RETURN VALUE**

38332 If successful, the *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()* functions shall return the  
38333 appropriate maximum or minimum values, respectively. If unsuccessful, they shall return a  
38334 value of  $-1$  and set *errno* to indicate the error.

38335 **ERRORS**

38336 The *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()* functions shall fail if:

38337 [EINVAL] The value of the *policy* parameter does not represent a defined scheduling  
38338 policy.

38339 **EXAMPLES**

38340 None.

38341 **APPLICATION USAGE**

38342 None.

38343 **RATIONALE**

38344 None.

38345 **FUTURE DIRECTIONS**

38346 None.

38347 **SEE ALSO**

38348 *sched\_getparam()*, *sched\_setparam()*, *sched\_getscheduler()*, *sched\_rr\_get\_interval()*,  
38349 *sched\_setscheduler()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**sched.h**>

38350 **CHANGE HISTORY**

38351 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38352 **Issue 6**

38353 These functions are marked as part of the Process Scheduling option.

38354 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38355 implementation does not support the Process Scheduling option.

38356 The [ESRCH] error condition has been removed since these functions do not take a *pid*  
38357 argument.

38358 **NAME**38359 sched\_getparam — get scheduling parameters (**REALTIME**)38360 **SYNOPSIS**

38361 PS #include &lt;sched.h&gt;

38362 int sched\_getparam(pid\_t pid, struct sched\_param \*param);

38363

38364 **DESCRIPTION**38365 The *sched\_getparam()* function shall return the scheduling parameters of a process specified by  
38366 *pid* in the **sched\_param** structure pointed to by *param*.38367 If a process specified by *pid* exists, and if the calling process has permission, the scheduling  
38368 parameters for the process whose process ID is equal to *pid* shall be returned.38369 If *pid* is zero, the scheduling parameters for the calling process shall be returned. The behavior of  
38370 the *sched\_getparam()* function is unspecified if the value of *pid* is negative.38371 **RETURN VALUE**38372 Upon successful completion, the *sched\_getparam()* function shall return zero. If the call to  
38373 *sched\_getparam()* is unsuccessful, the function shall return a value of -1 and set *errno* to indicate  
38374 the error.38375 **ERRORS**38376 The *sched\_getparam()* function shall fail if:38377 [EPERM] The requesting process does not have permission to obtain the scheduling |  
38378 parameters of the specified process.38379 [ESRCH] No process can be found corresponding to that specified by *pid*. |38380 **EXAMPLES**

38381 None.

38382 **APPLICATION USAGE**

38383 None.

38384 **RATIONALE**

38385 None.

38386 **FUTURE DIRECTIONS**

38387 None.

38388 **SEE ALSO**38389 *sched\_getscheduler()*, *sched\_setparam()*, *sched\_setscheduler()*, the Base Definitions volume of |  
38390 IEEE Std. 1003.1-200x, <**sched.h**>38391 **CHANGE HISTORY**

38392 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38393 **Issue 6**38394 The *sched\_getparam()* function is marked as part of the Process Scheduling option. |38395 The [ENOSYS] error condition has been removed as stubs need not be provided if an |  
38396 implementation does not support the Process Scheduling option.

38397 **NAME**

38398 sched\_getscheduler — get scheduling policy (**REALTIME**)

38399 **SYNOPSIS**

38400 PS #include <sched.h>

38401 int sched\_getscheduler(pid\_t pid);

38402

38403 **DESCRIPTION**

38404 The *sched\_getscheduler()* function shall return the scheduling policy of the process specified by  
 38405 *pid*. If the value of *pid* is negative, the behavior of the *sched\_getscheduler()* function is  
 38406 unspecified.

38407 The values that can be returned by *sched\_getscheduler()* are defined in the header file <**sched.h**>.

38408 If a process specified by *pid* exists, and if the calling process has permission, the scheduling  
 38409 policy shall be returned for the process whose process ID is equal to *pid*.

38410 If *pid* is zero, the scheduling policy shall be returned for the calling process.

38411 **RETURN VALUE**

38412 Upon successful completion, the *sched\_getscheduler()* function shall return the scheduling policy  
 38413 of the specified process. If unsuccessful, the function shall return  $-1$  and set *errno* to indicate the  
 38414 error.

38415 **ERRORS**

38416 The *sched\_getscheduler()* function shall fail if:

38417 [EPERM] The requesting process does not have permission to determine the scheduling  
 38418 policy of the specified process.

38419 [ESRCH] No process can be found corresponding to that specified by *pid*.

38420 **EXAMPLES**

38421 None.

38422 **APPLICATION USAGE**

38423 None.

38424 **RATIONALE**

38425 None.

38426 **FUTURE DIRECTIONS**

38427 None.

38428 **SEE ALSO**

38429 *sched\_getparam()*, *sched\_setparam()*, *sched\_setscheduler()*, the Base Definitions volume of  
 38430 IEEE Std. 1003.1-200x, <**sched.h**>

38431 **CHANGE HISTORY**

38432 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38433 **Issue 6**

38434 The *sched\_getscheduler()* function is marked as part of the Process Scheduling option.

38435 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 38436 implementation does not support the Process Scheduling option.

38437 **NAME**38438 sched\_rr\_get\_interval — get execution time limits (**REALTIME**)38439 **SYNOPSIS**

38440 PS #include &lt;sched.h&gt;

38441 int sched\_rr\_get\_interval(pid\_t *pid*, struct timespec \**interval*);

38442

38443 **DESCRIPTION**

38444 The *sched\_rr\_get\_interval()* function updates the **timespec** structure referenced by the *interval*  
 38445 argument to contain the current execution time limit (that is, time quantum) for the process  
 38446 specified by *pid*. If *pid* is zero, the current execution time limit for the calling process shall be  
 38447 returned.

38448 **RETURN VALUE**

38449 If successful, the *sched\_rr\_get\_interval()* function shall return zero. Otherwise, it shall return a  
 38450 value of  $-1$  and set *errno* to indicate the error.

38451 **ERRORS**38452 The *sched\_rr\_get\_interval()* function shall fail if:38453 [ESRCH] No process can be found corresponding to that specified by *pid*.38454 **EXAMPLES**

38455 None.

38456 **APPLICATION USAGE**

38457 None.

38458 **RATIONALE**

38459 None.

38460 **FUTURE DIRECTIONS**

38461 None.

38462 **SEE ALSO**

38463 *sched\_getparam()*, *sched\_get\_priority\_max()*, *sched\_getscheduler()*, *sched\_setparam()*,  
 38464 *sched\_setscheduler()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**sched.h**>

38465 **CHANGE HISTORY**

38466 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38467 **Issue 6**38468 The *sched\_rr\_get\_interval()* function is marked as part of the Process Scheduling option.

38469 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 38470 implementation does not support the Process Scheduling option.

38471 **NAME**

38472 sched\_setparam — set scheduling parameters (**REALTIME**)

38473 **SYNOPSIS**

38474 PS #include <sched.h>

38475 int sched\_setparam(pid\_t pid, const struct sched\_param \*param);

38476

38477 **DESCRIPTION**

38478 The *sched\_setparam()* function sets the scheduling parameters of the process specified by *pid* to the values specified by the **sched\_param** structure pointed to by *param*. The value of the *sched\_priority* member in the **sched\_param** structure is any integer within the inclusive priority range for the current scheduling policy of the process specified by *pid*. Higher numerical values for the priority represent higher priorities. If the value of *pid* is negative, the behavior of the *sched\_setparam()* function is unspecified.

38484 If a process specified by *pid* exists, and if the calling process has permission, the scheduling parameters shall be set for the process whose process ID is equal to *pid*.

38486 If *pid* is zero, the scheduling parameters shall be set for the calling process.

38487 The conditions under which one process has permission to change the scheduling parameters of another process are implementation-defined.

38489 Implementations may require the requesting process to have the appropriate privilege to set its own scheduling parameters or those of another process.

38491 The target process, whether it is running or not running, resumes execution after all other runnable processes of equal or greater priority have been scheduled to run.

38493 If the priority of the process specified by the *pid* argument is set higher than that of the lowest priority running process and if the specified process is ready to run, the process specified by the *pid* argument preempts a lowest priority running process. Similarly, if the process calling *sched\_setparam()* sets its own priority lower than that of one or more other non-empty process lists, then the process that is the head of the highest priority list also preempts the calling process. Thus, in either case, the originating process might not receive notification of the completion of the requested priority change until the higher priority process has executed.

38500 ss If the scheduling policy of the target process is SCHED\_SPORADIC, the value specified by the *sched\_ss\_low\_priority* member of the *param* argument shall be any integer within the inclusive priority range for the sporadic server policy. The *sched\_ss\_repl\_period* and *sched\_ss\_init\_budget* members of the *param* argument shall represent the time parameters to be used by the sporadic server scheduling policy for the target process. The *sched\_ss\_max\_repl* member of the *param* argument shall represent the maximum number of replenishments that are allowed to be pending simultaneously for the process scheduled under this scheduling policy.

38507 The specified *sched\_ss\_repl\_period* shall be greater than or equal to the specified *sched\_ss\_init\_budget* for the function to succeed; if it is not, then the function shall fail.

38509 The value of *sched\_ss\_max\_repl* shall be within the inclusive range [1,{SS\_REPL\_MAX}] for the function to succeed; if not, the function shall fail.

38511 If the scheduling policy of the target process is either SCHED\_FIFO or SCHED\_RR, the *sched\_ss\_low\_priority*, *sched\_ss\_repl\_period*, and *sched\_ss\_init\_budget* members of the *param* argument shall have no effect on the scheduling behavior. If the scheduling policy of this process is not SCHED\_FIFO, SCHED\_RR, or SCHED\_SPORADIC, including SCHED\_OTHER, the effects of these members shall be implementation-defined.

38516 **Notes to Reviewers**38517 *This section with side shading will not appear in the final copy. - Ed.*

38518 D3, XSH, ERN 502 questions "including SCHED\_OTHER" above. (Problem left over from D2.)

38519 If the current scheduling policy for the process specified by *pid* is not SCHED\_FIFO,  
38520 ss SCHED\_RR, or SCHED\_SPORADIC, the result is implementation-defined; this case includes the  
38521 SCHED\_OTHER policy.38522 The effect of this function on individual threads is dependent on the scheduling contention  
38523 scope of the threads:38524 • For threads with system scheduling contention scope, these functions have no effect on their  
38525 scheduling.38526 • For threads with process scheduling contention scope, the threads' scheduling parameters  
38527 shall not be affected. However, the scheduling of these threads with respect to threads in  
38528 other processes may be dependent on the scheduling parameters of their process, which are  
38529 governed using these functions.38530 If an implementation supports a two-level scheduling model in which library threads are  
38531 multiplexed on top of several kernel-scheduled entities, then the underlying kernel-scheduled  
38532 entities for the system contention scope threads shall not be affected by these functions.38533 The underlying kernel-scheduled entities for the process contention scope threads shall have  
38534 their scheduling parameters changed to the value specified in *param*. Kernel scheduled entities  
38535 for use by process contention scope threads that are created after this call completes inherit their  
38536 scheduling policy and associated scheduling parameters from the process.38537 This function is not atomic with respect to other threads in the process. Threads are allowed to  
38538 continue to execute while this function call is in the process of changing the scheduling policy  
38539 for the underlying kernel-scheduled entities used by the process contention scope threads.38540 **RETURN VALUE**38541 If successful, the *sched\_setparam()* function shall return zero.38542 If the call to *sched\_setparam()* is unsuccessful, the priority shall remain unchanged, and the  
38543 function shall return a value of -1 and set *errno* to indicate the error.38544 **ERRORS**38545 The *sched\_setparam()* function shall fail if:38546 [EINVAL] One or more of the requested scheduling parameters is outside the range  
38547 defined for the scheduling policy of the specified *pid*.38548 [EPERM] The requesting process does not have permission to set the scheduling  
38549 parameters for the specified process, or does not have the appropriate  
38550 privilege to invoke *sched\_setparam()*.38551 [ESRCH] No process can be found corresponding to that specified by *pid*.

38552 **EXAMPLES**

38553 None.

38554 **APPLICATION USAGE**

38555 None.

38556 **RATIONALE**

38557 None.

38558 **FUTURE DIRECTIONS**

38559 None.

38560 **SEE ALSO**

38561 *sched\_getparam()*, *sched\_getscheduler()*, *sched\_setscheduler()*, the Base Definitions volume of  
38562 IEEE Std. 1003.1-200x, <**sched.h**>

38563 **CHANGE HISTORY**

38564 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38565 **Issue 6**38566 The *sched\_setparam()* function is marked as part of the Process Scheduling option.

38567 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38568 implementation does not support the Process Scheduling option.

38569 The following new requirements on POSIX implementations derive from alignment with the  
38570 Single UNIX Specification:

38571 • In the DESCRIPTION, the effect of this function on a thread's scheduling parameters is  
38572 added.

38573 • Sections describing two-level scheduling and atomicity of the function are added.

38574 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std. 1003.1d-1999.

38575 **NAME**38576 sched\_setscheduler — set scheduling policy and parameters (**REALTIME**)38577 **SYNOPSIS**38578 PS `#include <sched.h>`38579 `int sched_setscheduler(pid_t pid, int policy,`  
38580 `const struct sched_param *param);`

38581

38582 **DESCRIPTION**

38583 The *sched\_setscheduler()* function sets the scheduling policy and scheduling parameters of the  
 38584 process specified by *pid* to *policy* and the parameters specified in the **sched\_param** structure  
 38585 pointed to by *param*, respectively. The value of the *sched\_priority* member in the **sched\_param**  
 38586 structure is any integer within the inclusive priority range for the scheduling policy specified by  
 38587 *policy*. If the value of *pid* is negative, the behavior of the *sched\_setscheduler()* function is  
 38588 unspecified.

38589 The possible values for the *policy* parameter are defined in the header file **<sched.h>**.

38590 If a process specified by *pid* exists, and if the calling process has permission, the scheduling  
 38591 policy and scheduling parameters shall be set for the process whose process ID is equal to *pid*.

38592 If *pid* is zero, the scheduling policy and scheduling parameters shall be set for the calling  
 38593 process.

38594 The conditions under which one process has the appropriate privilege to change the scheduling  
 38595 parameters of another process are implementation-defined.

38596 Implementations may require that the requesting process have permission to set its own  
 38597 scheduling parameters or those of another process. Additionally, implementation-defined  
 38598 restrictions may apply as to the appropriate privileges required to set a process' own scheduling  
 38599 policy, or another process' scheduling policy, to a particular value.

38600 The *sched\_setscheduler()* function is considered successful if it succeeds in setting the scheduling  
 38601 policy and scheduling parameters of the process specified by *pid* to the values specified by *policy*  
 38602 and the structure pointed to by *param*, respectively.

38603 ss If the scheduling policy specified by *policy* is **SCHED\_SPORADIC**, the value specified by the  
 38604 *sched\_ss\_low\_priority* member of the *param* argument shall be any integer within the inclusive  
 38605 priority range for the sporadic server policy. The *sched\_ss\_repl\_period* and *sched\_ss\_init\_budget*  
 38606 members of the *param* argument shall represent the time parameters used by the sporadic server  
 38607 scheduling policy for the target process. The *sched\_ss\_max\_repl* member of the *param* argument  
 38608 shall represent the maximum number of replenishments that are allowed to be pending  
 38609 simultaneously for the process scheduled under this scheduling policy.

38610 The specified *sched\_ss\_repl\_period* shall be greater than or equal to the specified  
 38611 *sched\_ss\_init\_budget* for the function to succeed; if it is not, then the function shall fail.

38612 The value of *sched\_ss\_max\_repl* shall be within the inclusive range [1,{**SS\_REPL\_MAX**}] for the  
 38613 function to succeed; if not, the function shall fail.

38614 If the scheduling policy specified by *policy* is either **SCHED\_FIFO** or **SCHED\_RR**, the  
 38615 *sched\_ss\_low\_priority*, *sched\_ss\_repl\_period*, and *sched\_ss\_init\_budget* members of the *param*  
 38616 argument shall have no effect on the scheduling behavior.

38617 The effect of this function on individual threads is dependent on the scheduling contention  
 38618 scope of the threads:

38619           • For threads with system scheduling contention scope, these functions have no effect on their  
38620 scheduling.

38621           • For threads with process scheduling contention scope, the threads' scheduling policy and  
38622 associated parameters shall not be affected. However, the scheduling of these threads with  
38623 respect to threads in other processes may be dependent on the scheduling parameters of their  
38624 process, which are governed using these functions.

38625           If an implementation supports a two-level scheduling model in which library threads are  
38626 multiplexed on top of several kernel-scheduled entities, then the underlying kernel-scheduled  
38627 entities for the system contention scope threads shall not be affected by these functions.

38628           The underlying kernel-scheduled entities for the process contention scope threads shall have  
38629 their scheduling policy and associated scheduling parameters changed to the values specified in  
38630 *policy* and *param*, respectively. Kernel scheduled entities for use by process contention scope  
38631 threads that are created after this call completes inherit their scheduling policy and associated  
38632 scheduling parameters from the process.

38633           This function is not atomic with respect to other threads in the process. Threads are allowed to  
38634 continue to execute while this function call is in the process of changing the scheduling policy  
38635 and associated scheduling parameters for the underlying kernel-scheduled entities used by the  
38636 process contention scope threads.

#### 38637 RETURN VALUE

38638           Upon successful completion, the function shall return the former scheduling policy of the  
38639 specified process. If the *sched\_setscheduler()* function fails to complete successfully, the policy  
38640 and scheduling parameters shall remain unchanged, and the function shall return a value of  $-1$   
38641 and set *errno* to indicate the error.

#### 38642 ERRORS

38643           The *sched\_setscheduler()* function shall fail if:

38644           [EINVAL]           The value of the *policy* parameter is invalid, or one or more of the parameters  
38645 contained in *param* is outside the valid range for the specified scheduling  
38646 policy.

38647           [EPERM]           The requesting process does not have permission to set either or both of the  
38648 scheduling parameters or the scheduling policy of the specified process.

38649           [ESRCH]           No process can be found corresponding to that specified by *pid*.

#### 38650 EXAMPLES

38651           None.

#### 38652 APPLICATION USAGE

38653           None.

#### 38654 RATIONALE

38655           None.

#### 38656 FUTURE DIRECTIONS

38657           None.

#### 38658 SEE ALSO

38659           *sched\_getparam()*, *sched\_getscheduler()*, *sched\_setparam()*, the Base Definitions volume of  
38660 IEEE Std. 1003.1-200x, <*sched.h*>

38661 **CHANGE HISTORY**

38662 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38663 **Issue 6**

38664 The *sched\_setscheduler()* function is marked as part of the Process Scheduling option.

38665 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38666 implementation does not support the Process Scheduling option.

38667 The following new requirements on POSIX implementations derive from alignment with the  
38668 Single UNIX Specification:

38669 • In the DESCRIPTION, the effect of this function on a thread's scheduling parameters is  
38670 added.

38671 • Sections describing two-level scheduling and atomicity of the function are added.

38672 The SCHED\_SPORADIC scheduling policy is added for alignment with IEEE Std. 1003.1d-1999.

38673 **NAME**

38674 sched\_yield — yield processor

38675 **SYNOPSIS**

38676 PS|THR #include &lt;sched.h&gt;

38677 int sched\_yield(void);

38678

38679 **DESCRIPTION**38680 The *sched\_yield()* function forces the running thread to relinquish the processor until it again  
38681 becomes the head of its thread list. It takes no arguments.38682 **RETURN VALUE**38683 The *sched\_yield()* function shall return 0 if it completes successfully; otherwise, it shall return a  
38684 value of -1 and set *errno* to indicate the error.38685 **ERRORS**

38686 No errors are defined.

38687 **EXAMPLES**

38688 None.

38689 **APPLICATION USAGE**

38690 None.

38691 **RATIONALE**

38692 None.

38693 **FUTURE DIRECTIONS**

38694 None.

38695 **SEE ALSO**

38696 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;sched.h&gt;

38697 **CHANGE HISTORY**38698 First released in Issue 5. Included for alignment with the POSIX Realtime Extension and the  
38699 POSIX Threads Extension.38700 **Issue 6**38701 The *sched\_yield()* function is now marked as part of the Process Scheduling and Threads options.

38702 **NAME**

38703        seed48 — seed uniformly distributed pseudo-random non-negative long integer generator

38704 **SYNOPSIS**

38705 xSI        #include &lt;stdlib.h&gt;

38706        unsigned short \*seed48(unsigned short *seed16v*[3]);

38707

38708 **DESCRIPTION**38709        Refer to *drand48()*.

38710 **NAME**

38711 seekdir — set position of directory stream

38712 **SYNOPSIS**

38713 XSI #include &lt;dirent.h&gt;

38714 void seekdir(DIR \*dirp, long loc);

38715

38716 **DESCRIPTION**

38717 The *seekdir()* function sets the position of the next *readdir()* operation on the directory stream  
38718 specified by *dirp* to the position specified by *loc*. The value of *loc* should have been returned  
38719 from an earlier call to *telldir()*. The new position reverts to the one associated with the directory  
38720 stream when *telldir()* was performed.

38721 If the value of *loc* was not obtained from an earlier call to *telldir()*, or if a call to *rewinddir()*  
38722 occurred between the call to *telldir()* and the call to *seekdir()*, the results of subsequent calls to  
38723 *readdir()* are unspecified.

38724 **RETURN VALUE**38725 The *seekdir()* function shall return no value.38726 **ERRORS**

38727 No errors are defined.

38728 **EXAMPLES**

38729 None.

38730 **APPLICATION USAGE**

38731 None.

38732 **RATIONALE**

38733 None.

38734 **FUTURE DIRECTIONS**

38735 None.

38736 **SEE ALSO**38737 *opendir()*, *readdir()*, *telldir()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <dirent.h>,

38738 &lt;stdio.h&gt;, &lt;sys/types.h&gt;

38739 **CHANGE HISTORY**

38740 First released in Issue 2.

38741 **Issue 4**

38742 The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
38743 XSI-conformant systems.

38744 The type of argument *loc* is expanded to **long**.38745 **Issue 4, Version 2**

38746 The DESCRIPTION is updated for X/OPEN UNIX conformance to indicate that a call to  
38747 *readdir()* may produce unspecified results if either *loc* was not obtained by a previous call to  
38748 *telldir()*, or if there is an intervening call to *rewinddir()*.

38749 **Issue 6**

38750 In the SYNOPSIS, the inclusion of &lt;sys/types.h&gt; is no longer required.

38751 **NAME**

38752       select — synchronous I/O multiplexing

38753 **SYNOPSIS**

38754       #include &lt;sys/time.h&gt;

38755       int select(int *nfds*, fd\_set \*restrict *readfds*, fd\_set \*restrict *writefds*,38756               fd\_set \*restrict *errorfds*, struct timeval \*restrict *timeout*);

38757

38758 **DESCRIPTION**38759       Refer to *pselect()*.

38760 **NAME**

38761 sem\_close — close a named semaphore (**REALTIME**)

38762 **SYNOPSIS**

38763 SEM #include <semaphore.h>

38764 int sem\_close(sem\_t \*sem);

38765

38766 **DESCRIPTION**

38767 The *sem\_close()* function is used to indicate that the calling process is finished using the named  
38768 semaphore indicated by *sem*. The effects of calling *sem\_close()* for an unnamed semaphore (one  
38769 created by *sem\_init()*) are undefined. The *sem\_close()* function deallocates (that is, makes  
38770 available for reuse by a subsequent *sem\_open()* by this process) any system resources allocated  
38771 by the system for use by this process for this semaphore. The effect of subsequent use of the  
38772 semaphore indicated by *sem* by this process is undefined. If the semaphore has not been  
38773 removed with a successful call to *sem\_unlink()*, then *sem\_close()* has no effect on the state of the  
38774 semaphore. If the *sem\_unlink()* function has been successfully invoked for *name* after the most  
38775 recent call to *sem\_open()* with *O\_CREAT* for this semaphore, then when all processes that have  
38776 opened the semaphore close it, the semaphore is no longer accessible.

38777 **RETURN VALUE**

38778 Upon successful completion, a value of zero shall be returned. Otherwise, a value of -1 shall be  
38779 returned and *errno* set to indicate the error.

38780 **ERRORS**

38781 The *sem\_close()* function shall fail if:

38782 [EINVAL] The *sem* argument is not a valid semaphore descriptor.

38783 **EXAMPLES**

38784 None.

38785 **APPLICATION USAGE**

38786 The *sem\_close()* function is part of the Semaphores option and need not be available on all  
38787 implementations.

38788 **RATIONALE**

38789 None.

38790 **FUTURE DIRECTIONS**

38791 None.

38792 **SEE ALSO**

38793 *semctl()*, *semget()*, *semop()*, *sem\_init()*, *sem\_open()*, *sem\_unlink()*, the Base Definitions volume of  
38794 IEEE Std. 1003.1-200x, <semaphore.h>

38795 **CHANGE HISTORY**

38796 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38797 **Issue 6**

38798 The *sem\_close()* function is marked as part of the Semaphores option.

38799 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38800 implementation does not support the Semaphores option.

38801 **NAME**38802 sem\_destroy — destroy an unnamed semaphore (**REALTIME**)38803 **SYNOPSIS**

38804 SEM #include &lt;semaphore.h&gt;

38805 int sem\_destroy(sem\_t \*sem);

38806

38807 **DESCRIPTION**

38808 The *sem\_destroy()* function is used to destroy the unnamed semaphore indicated by *sem*. Only a  
38809 semaphore that was created using *sem\_init()* may be destroyed using *sem\_destroy()*; the effect of  
38810 calling *sem\_destroy()* with a named semaphore is undefined. The effect of subsequent use of the  
38811 semaphore *sem* is undefined until *sem* is re-initialized by another call to *sem\_init()*.

38812 It is safe to destroy an initialized semaphore upon which no threads are currently blocked. The  
38813 effect of destroying a semaphore upon which other threads are currently blocked is undefined.

38814 **RETURN VALUE**

38815 Upon successful completion, a value of zero shall be returned. Otherwise, a value of -1 shall be  
38816 returned and *errno* set to indicate the error.

38817 **ERRORS**38818 The *sem\_destroy()* function shall fail if:38819 [EINVAL] The *sem* argument is not a valid semaphore.38820 The *sem\_destroy()* function may fail if:

38821 [EBUSY] There are currently processes blocked on the semaphore.

38822 **EXAMPLES**

38823 None.

38824 **APPLICATION USAGE**

38825 The *sem\_destroy()* function is part of the Semaphores option and need not be available on all  
38826 implementations.

38827 **RATIONALE**

38828 None.

38829 **FUTURE DIRECTIONS**

38830 None.

38831 **SEE ALSO**

38832 *semctl()*, *semget()*, *semop()*, *sem\_init()*, *sem\_open()*, the Base Definitions volume of  
38833 IEEE Std. 1003.1-200x, <semaphore.h>

38834 **CHANGE HISTORY**

38835 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38836 **Issue 6**38837 The *sem\_destroy()* function is marked as part of the Semaphores option.

38838 The [ENOSYS] error condition has been removed as stubs need not be provided if an

38839 implementation does not support the Semaphores option.

38840 **NAME**

38841 sem\_getvalue — get the value of a semaphore (**REALTIME**)

38842 **SYNOPSIS**

```
38843 SEM #include <semaphore.h>
```

```
38844 int sem_getvalue(sem_t *restrict sem, int *restrict sval);
```

38845

38846 **DESCRIPTION**

38847 The *sem\_getvalue()* function updates the location referenced by the *sval* argument to have the  
38848 value of the semaphore referenced by *sem* without affecting the state of the semaphore. The  
38849 updated value represents an actual semaphore value that occurred at some unspecified time  
38850 during the call, but it need not be the actual value of the semaphore when it is returned to the  
38851 calling process.

38852 If *sem* is locked, then the value returned by *sem\_getvalue()* is either zero or a negative number  
38853 whose absolute value represents the number of processes waiting for the semaphore at some  
38854 unspecified time during the call.

38855 **RETURN VALUE**

38856 Upon successful completion, the *sem\_getvalue()* function shall return a value of zero. Otherwise,  
38857 it shall return a value of -1 and set *errno* to indicate the error.

38858 **ERRORS**

38859 The *sem\_getvalue()* function shall fail if:

38860 [EINVAL] The *sem* argument does not refer to a valid semaphore.

38861 **EXAMPLES**

38862 None.

38863 **APPLICATION USAGE**

38864 The *sem\_getvalue()* function is part of the Semaphores option and need not be available on all  
38865 implementations.

38866 **RATIONALE**

38867 None.

38868 **FUTURE DIRECTIONS**

38869 None.

38870 **SEE ALSO**

38871 *semctl()*, *semget()*, *semop()*, *sem\_post()*, *sem\_timedwait()*, *sem\_trywait()*, *sem\_wait()*, the Base  
38872 Definitions volume of IEEE Std. 1003.1-200x, <**semaphore.h**>

38873 **CHANGE HISTORY**

38874 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38875 **Issue 6**

38876 The *sem\_getvalue()* function is marked as part of the Semaphores option.

38877 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38878 implementation does not support the Semaphores option.

38879 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
38880 IEEE Std. 1003.1d-1999.

38881 The **restrict** keyword is added to the *sem\_getvalue()* prototype for alignment with the  
38882 ISO/IEC 9899:1999 standard.

38883 **NAME**38884 sem\_init — initialize an unnamed semaphore (**REALTIME**)38885 **SYNOPSIS**

38886 SEM #include &lt;semaphore.h&gt;

38887 int sem\_init(sem\_t \*sem, int pshared, unsigned value);

38888

38889 **DESCRIPTION**

38890 The *sem\_init()* function is used to initialize the unnamed semaphore referred to by *sem*. The  
 38891 value of the initialized semaphore is *value*. Following a successful call to *sem\_init()*, the  
 38892 semaphore may be used in subsequent calls to *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and  
 38893 *sem\_destroy()*. This semaphore remains usable until the semaphore is destroyed.

38894 If the *pshared* argument has a non-zero value, then the semaphore is shared between processes;  
 38895 in this case, any process that can access the semaphore *sem* can use *sem* for performing  
 38896 *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and *sem\_destroy()* operations.

38897 Only *sem* itself may be used for performing synchronization. The result of referring to copies of  
 38898 *sem* in calls to *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and *sem\_destroy()*, is undefined.

38899 If the *pshared* argument is zero, then the semaphore is shared between threads of the process; any  
 38900 thread in this process can use *sem* for performing *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and  
 38901 *sem\_destroy()* operations. The use of the semaphore by threads other than those created in the  
 38902 same process is undefined.

38903 Attempting to initialize an already initialized semaphore results in undefined behavior.

38904 **RETURN VALUE**

38905 Upon successful completion, the *sem\_init()* function shall initialize the semaphore in *sem*.  
 38906 Otherwise, it shall return  $-1$  and set *errno* to indicate the error.

38907 **ERRORS**

38908 The *sem\_init()* function shall fail if:

38909 [EINVAL] The *value* argument exceeds {SEM\_VALUE\_MAX}.

38910 [ENOSPC] A resource required to initialize the semaphore has been exhausted, or the  
 38911 limit on semaphores ({SEM\_NSEMS\_MAX}) has been reached.

38912 [EPERM] The process lacks the appropriate privileges to initialize the semaphore.

38913 **EXAMPLES**

38914 None.

38915 **APPLICATION USAGE**

38916 The *sem\_init()* function is part of the Semaphores option and need not be available on all  
 38917 implementations.

38918 **RATIONALE**

38919 Although this volume of IEEE Std. 1003.1-200x fails to specify a successful return value, it is  
 38920 likely that a later version may require the implementation to return a value of zero if the call to  
 38921 *sem\_init()* is successful.

38922 **FUTURE DIRECTIONS**

38923 None.

38924 **SEE ALSO**

38925 *sem\_destroy()*, *sem\_post()*, *sem\_timedwait()*, *sem\_trywait()*, *sem\_wait()*, the Base Definitions  
38926 volume of IEEE Std. 1003.1-200x, <semaphore.h>

38927 **CHANGE HISTORY**

38928 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

38929 **Issue 6**

38930 The *sem\_init()* function is marked as part of the Semaphores option.

38931 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
38932 implementation does not support the Semaphores option.

38933 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
38934 IEEE Std. 1003.1d-1999.

38935 **NAME**38936 sem\_open — initialize and open a named semaphore (**REALTIME**)38937 **SYNOPSIS**

38938 SEM #include &lt;semaphore.h&gt;

38939 sem\_t \*sem\_open(const char \*name, int oflag, ...);

38940

38941 **DESCRIPTION**

38942 The *sem\_open()* function establishes a connection between a named semaphore and a process.  
 38943 Following a call to *sem\_open()* with semaphore name *name*, the process may reference the  
 38944 semaphore associated with *name* using the address returned from the call. This semaphore may  
 38945 be used in subsequent calls to *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and *sem\_close()*. The  
 38946 semaphore remains usable by this process until the semaphore is closed by a successful call to  
 38947 *sem\_close()*, *\_exit()*, or one of the *exec* functions.

38948 The *oflag* argument controls whether the semaphore is created or merely accessed by the call to  
 38949 *sem\_open()*. The following flag bits may be set in *oflag*:

38950 **O\_CREAT** This flag is used to create a semaphore if it does not already exist. If **O\_CREAT**  
 38951 is set and the semaphore already exists, then **O\_CREAT** has no effect, except as noted  
 38952 under **O\_EXCL**. Otherwise, *sem\_open()* creates a named semaphore. The **O\_CREAT**  
 38953 flag requires a third and a fourth argument: *mode*, which is of type **mode\_t**, and  
 38954 *value*, which is of type **unsigned**. The semaphore is created with an initial value of  
 38955 *value*. Valid initial values for semaphores are less than or equal to  
 38956 {SEM\_VALUE\_MAX}.

38957 The user ID of the semaphore is set to the effective user ID of the process; the  
 38958 group ID of the semaphore is set to a system default group ID or to the effective  
 38959 group ID of the process. The permission bits of the semaphore are set to the value  
 38960 of the *mode* argument except those set in the file mode creation mask of the  
 38961 process. When bits in *mode* other than the file permission bits are specified, the  
 38962 effect is unspecified.

38963 After the semaphore named *name* has been created by *sem\_open()* with the  
 38964 **O\_CREAT** flag, other processes can connect to the semaphore by calling  
 38965 *sem\_open()* with the same value of *name*.

38966 **O\_EXCL** If **O\_EXCL** and **O\_CREAT** are set, *sem\_open()* fails if the semaphore *name* exists.  
 38967 The check for the existence of the semaphore and the creation of the semaphore if  
 38968 it does not exist are atomic with respect to other processes executing *sem\_open()*  
 38969 with **O\_EXCL** and **O\_CREAT** set. If **O\_EXCL** is set and **O\_CREAT** is not set, the  
 38970 effect is undefined.

38971 If flags other than **O\_CREAT** and **O\_EXCL** are specified in the *oflag* parameter, the  
 38972 effect is unspecified.

38973 The *name* argument points to a string naming a semaphore object. It is unspecified whether the  
 38974 name appears in the file system and is visible to functions that take path names as arguments.  
 38975 The *name* argument conforms to the construction rules for a path name. If *name* begins with the  
 38976 slash character, then processes calling *sem\_open()* with the same value of *name* shall refer to the  
 38977 same semaphore object, as long as that name has not been removed. If *name* does not begin with  
 38978 the slash character, the effect is implementation-defined. The interpretation of slash characters  
 38979 other than the leading slash character in *name* is implementation-defined.

38980 If a process makes multiple successful calls to *sem\_open()* with the same value for *name*, the  
 38981 same semaphore address is returned for each such successful call, provided that there have been

38982 no calls to *sem\_unlink()* for this semaphore.

38983 References to copies of the semaphore produce undefined results.

#### 38984 RETURN VALUE

38985 Upon successful completion, the *sem\_open()* function shall return the address of the semaphore.  
 38986 Otherwise, it shall return a value of SEM\_FAILED and set *errno* to indicate the error. The symbol  
 38987 SEM\_FAILED is defined in the header <semaphore.h>. No successful return from *sem\_open()*  
 38988 shall return the value SEM\_FAILED.

#### 38989 ERRORS

38990 If any of the following conditions occur, the *sem\_open()* function shall return SEM\_FAILED and  
 38991 set *errno* to the corresponding value:

38992 [EACCES] The named semaphore exists and the permissions specified by *oflag* are  
 38993 denied, or the named semaphore does not exist and permission to create the  
 38994 named semaphore is denied.

38995 [EEXIST] O\_CREAT and O\_EXCL are set and the named semaphore already exists.

38996 [EINTR] The *sem\_open()* operation was interrupted by a signal.

38997 [EINVAL] The *sem\_open()* operation is not supported for the given name, or O\_CREAT  
 38998 was specified in *oflag* and *value* was greater than {SEM\_VALUE\_MAX}.

38999 [EMFILE] Too many semaphore descriptors or file descriptors are currently in use by  
 39000 this process.

39001 [ENAMETOOLONG]

39002 The length of the *name* argument exceeds {PATH\_MAX} or a path name  
 39003 component is longer than {NAME\_MAX}.

39004 [ENFILE] Too many semaphores are currently open in the system.

39005 [ENOENT] O\_CREAT is not set and the named semaphore does not exist.

39006 [ENOSPC] There is insufficient space for the creation of the new named semaphore.

#### 39007 EXAMPLES

39008 None.

#### 39009 APPLICATION USAGE

39010 The *sem\_open()* function is part of the Semaphores option and need not be available on all  
 39011 implementations.

#### 39012 RATIONALE

39013 An earlier version of this volume of IEEE Std. 1003.1-200x required an error return value of -1  
 39014 with the type **sem\_t\*** for the *sem\_open()* function, which is not guaranteed to be portable across  
 39015 implementations. The revised text provides the symbolic error code SEM\_FAILED to eliminate  
 39016 the type conflict.

#### 39017 FUTURE DIRECTIONS

39018 None.

#### 39019 SEE ALSO

39020 *semctl()*, *semget()*, *semop()*, *sem\_close()*, *sem\_post()*, *sem\_timedwait()*, *sem\_trywait()*, *sem\_unlink()*,  
 39021 *sem\_wait()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <semaphore.h>

39022 **CHANGE HISTORY**

39023 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39024 **Issue 6**

39025 The *sem\_open()* function is marked as part of the Semaphores option.

39026 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39027 implementation does not support the Semaphores option.

39028 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
39029 IEEE Std. 1003.1d-1999.

## 39030 NAME

39031 sem\_post — unlock a semaphore (**REALTIME**)

## 39032 SYNOPSIS

39033 SEM #include <semaphore.h>

39034 int sem\_post(sem\_t \*sem);

39035

## 39036 DESCRIPTION

39037 The *sem\_post()* function unlocks the semaphore referenced by *sem* by performing a semaphore  
39038 unlock operation on that semaphore.

39039 If the semaphore value resulting from this operation is positive, then no threads were blocked  
39040 waiting for the semaphore to become unlocked; the semaphore value is simply incremented.

39041 If the value of the semaphore resulting from this operation is zero, then one of the threads  
39042 blocked waiting for the semaphore shall be allowed to return successfully from its call to  
39043 *sem\_wait()*. If the Process Scheduling option is supported, the thread to be unblocked shall be  
39044 chosen in a manner appropriate to the scheduling policies and parameters in effect for the  
39045 blocked threads. In the case of the schedulers SCHED\_FIFO and SCHED\_RR, the highest  
39046 priority waiting thread shall be unblocked, and if there is more than one highest priority thread  
39047 blocked waiting for the semaphore, then the highest priority thread that has been waiting the  
39048 longest shall be unblocked. If the Process Scheduling option is not defined, the choice of a thread  
39049 to unblock is unspecified.

39050 SS If the Process Sporadic Server option is supported, and the scheduling policy is  
39051 SCHED\_SPORADIC, the semantics are as per SCHED\_FIFO above.

39052 The *sem\_post()* function shall be reentrant with respect to signals and may be invoked from a  
39053 signal-catching function.

## 39054 RETURN VALUE

39055 If successful, the *sem\_post()* function shall return zero; otherwise, the function shall return -1  
39056 and set *errno* to indicate the error.

## 39057 ERRORS

39058 The *sem\_post()* function shall fail if:

39059 [EINVAL] The *sem* argument does not refer to a valid semaphore.

## 39060 EXAMPLES

39061 None.

## 39062 APPLICATION USAGE

39063 The *sem\_post()* function is part of the Semaphores option and need not be available on all  
39064 implementations.

## 39065 RATIONALE

39066 None.

## 39067 FUTURE DIRECTIONS

39068 None.

## 39069 SEE ALSO

39070 *semctl()*, *semget()*, *semop()*, *sem\_timedwait()*, *sem\_trywait()*, *sem\_wait()*, the Base Definitions  
39071 volume of IEEE Std. 1003.1-200x, <semaphore.h>

39072 **CHANGE HISTORY**

39073 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39074 **Issue 6**

39075 The *sem\_post()* function is marked as part of the Semaphores option.

39076 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39077 implementation does not support the Semaphores option.

39078 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
39079 IEEE Std. 1003.1d-1999.

39080 SCHED\_SPORADIC is added to the list of scheduling policies for which the thread that is to be  
39081 unblocked is specified for alignment with IEEE Std. 1003.1d-1999.

## 39082 NAME

39083 sem\_timedwait — lock a semaphore (**REALTIME**)

## 39084 SYNOPSIS

39085 SEM TMO #include &lt;semaphore.h&gt;

39086 #include &lt;time.h&gt;

39087 int sem\_timedwait(sem\_t \*restrict sem,

39088 const struct timespec \*restrict abs\_timeout);

39089

## 39090 DESCRIPTION

39091 The *sem\_timedwait()* function locks the semaphore referenced by *sem* as in the *sem\_wait()*  
 39092 function. However, if the semaphore cannot be locked without waiting for another process or  
 39093 thread to unlock the semaphore by performing a *sem\_post()* function, this wait shall be  
 39094 terminated when the specified timeout expires.

39095 The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the  
 39096 clock on which timeouts are based (that is, when the value of that clock equals or exceeds  
 39097 *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time  
 39098 of the call. If the Timers option is supported, the timeout is based on the CLOCK\_REALTIME  
 39099 clock; if the Timers option is not supported, the timeout is based on the system clock as returned  
 39100 by the *time()* function.

## 39101 RETURN VALUE

39102 The *sem\_timedwait()* function shall return zero if the calling process successfully performed the  
 39103 semaphore lock operation on the semaphore designated by *sem*. If the call was unsuccessful, the  
 39104 state of the semaphore shall be unchanged, and the function shall return a value of -1 and set  
 39105 *errno* to indicate the error.

## 39106 ERRORS

39107 The *sem\_timedwait()* function shall fail if:39108 [EINVAL] The *sem* argument does not refer to a valid semaphore.

39109 [EINVAL] The process or thread would have blocked, and the *abs\_timeout* parameter  
 39110 specified a nanoseconds field value less than zero or greater than or equal to  
 39111 1 000 million.

39112 [ETIMEDOUT] The semaphore could not be locked before the specified timeout expired.

39113 The *sem\_timedwait()* function may fail if:

39114 [EDEADLK] A deadlock condition was detected.

39115 [EINTR] A signal interrupted this function.

## 39116 EXAMPLES

39117 None.

## 39118 APPLICATION USAGE

39119 Applications using these functions may be subject to priority inversion, as discussed in the Base  
 39120 Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

39121 The *sem\_timedwait()* function is part of the Semaphores and Timeouts options and need not be  
 39122 provided on all implementations.

39123 **RATIONALE**

39124 None.

39125 **FUTURE DIRECTIONS**

39126 None.

39127 **SEE ALSO**39128 *sem\_post()*, *sem\_trywait()*, *sem\_wait()*, *semctl()*, *semget()*, *semop()*, *time()*, the Base Definitions |39129 volume of IEEE Std. 1003.1-200x, <**semaphore.h**>, <**time.h**> |39130 **CHANGE HISTORY**

39131 First released in Issue 6. Derived from IEEE Std. 1003.1d-1999.

39132 **NAME**39133 sem\_trywait, sem\_wait — lock a semaphore (**REALTIME**)39134 **SYNOPSIS**

39135 SEM #include &lt;semaphore.h&gt;

39136 int sem\_trywait(sem\_t \*sem);

39137 int sem\_wait(sem\_t \*sem);

39138

39139 **DESCRIPTION**

39140 The *sem\_trywait()* function locks the semaphore referenced by *sem* only if the semaphore is  
 39141 currently not locked; that is, if the semaphore value is currently positive. Otherwise, it does not  
 39142 lock the semaphore.

39143 The *sem\_wait()* function locks the semaphore referenced by *sem* by performing a semaphore lock  
 39144 operation on that semaphore. If the semaphore value is currently zero, then the calling thread  
 39145 shall not return from the call to *sem\_wait()* until it either locks the semaphore or the call is  
 39146 interrupted by a signal.

39147 Upon successful return, the state of the semaphore shall be locked and shall remain locked until  
 39148 the *sem\_post()* function is executed and returns successfully.

39149 The *sem\_wait()* function is interruptible by the delivery of a signal.

39150 **RETURN VALUE**

39151 The *sem\_trywait()* and *sem\_wait()* functions shall return zero if the calling process successfully  
 39152 performed the semaphore lock operation on the semaphore designated by *sem*. If the call was  
 39153 unsuccessful, the state of the semaphore shall be unchanged, and the function shall return a  
 39154 value of  $-1$  and set *errno* to indicate the error.

39155 **ERRORS**

39156 The *sem\_trywait()* and *sem\_wait()* functions shall fail if:

39157 [EAGAIN] The semaphore was already locked, so it cannot be immediately locked by the  
 39158 *sem\_trywait()* operation (*sem\_trywait()* only).

39159 [EINVAL] The *sem* argument does not refer to a valid semaphore.

39160 The *sem\_trywait()* and *sem\_wait()* functions may fail if:

39161 [EDEADLK] A deadlock condition was detected.

39162 [EINTR] A signal interrupted this function.

39163 **EXAMPLES**

39164 None.

39165 **APPLICATION USAGE**

39166 Applications using these functions may be subject to priority inversion, as discussed in the Base  
 39167 Definitions volume of IEEE Std. 1003.1-200x, Section 3.287, Priority Inversion.

39168 The *sem\_trywait()* and *sem\_wait()* functions are part of the Semaphores option and need not be  
 39169 provided on all implementations.

39170 **RATIONALE**

39171 None.

39172 **FUTURE DIRECTIONS**

39173 None.

39174 **SEE ALSO**

39175 *semctl()*, *semget()*, *semop()*, *sem\_post()*, *sem\_timedwait()*, the Base Definitions volume of  
39176 IEEE Std. 1003.1-200x, <**semaphore.h**>

39177 **CHANGE HISTORY**

39178 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39179 **Issue 6**39180 The *sem\_trywait()* and *sem\_wait()* functions are marked as part of the Semaphores option.

39181 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
39182 implementation does not support the Semaphores option.

39183 The *sem\_timedwait()* function is added to the SEE ALSO section for alignment with  
39184 IEEE Std. 1003.1d-1999.

39185 **NAME**39186 sem\_unlink — remove a named semaphore (**REALTIME**)39187 **SYNOPSIS**

39188 SEM #include &lt;semaphore.h&gt;

39189 int sem\_unlink(const char \*name);

39190

39191 **DESCRIPTION**

39192 The *sem\_unlink()* function removes the semaphore named by the string *name*. If the semaphore  
 39193 named by *name* is currently referenced by other processes, then *sem\_unlink()* has no effect on the  
 39194 state of the semaphore. If one or more processes have the semaphore open when *sem\_unlink()* is  
 39195 called, destruction of the semaphore is postponed until all references to the semaphore have  
 39196 been destroyed by calls to *sem\_close()*, *\_exit()*, or *exec*. Calls to *sem\_open()* to recreate or  
 39197 reconnect to the semaphore refer to a new semaphore after *sem\_unlink()* is called. The  
 39198 *sem\_unlink()* call does not block until all references have been destroyed; it shall return  
 39199 immediately.

39200 **RETURN VALUE**

39201 Upon successful completion, the *sem\_unlink()* function shall return a value of 0. Otherwise, the  
 39202 semaphore shall not be changed and the function shall return a value of -1 and set *errno* to  
 39203 indicate the error.

39204 **ERRORS**39205 The *sem\_unlink()* function shall fail if:

39206 [EACCES] Permission is denied to unlink the named semaphore.

39207 [ENAMETOOLONG]

39208 The length of the *name* argument exceeds {PATH\_MAX} or a path name  
 39209 component is longer than {NAME\_MAX}.

39210 [ENOENT] The named semaphore does not exist.

39211 **EXAMPLES**

39212 None.

39213 **APPLICATION USAGE**

39214 The *sem\_unlink()* function is part of the Semaphores option and need not be available on all  
 39215 implementations.

39216 **RATIONALE**

39217 None.

39218 **FUTURE DIRECTIONS**

39219 None.

39220 **SEE ALSO**

39221 *semctl()*, *semget()*, *semop()*, *sem\_close()*, *sem\_open()*, the Base Definitions volume of  
 39222 IEEE Std. 1003.1-200x, <semaphore.h>

39223 **CHANGE HISTORY**

39224 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

39225 **Issue 6**39226 The *sem\_unlink()* function is marked as part of the Semaphores option.

39227 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
 39228 implementation does not support the Semaphores option.



39229 **NAME**

39230 sem\_wait — lock a semaphore (**REALTIME**)

39231 **SYNOPSIS**

39232 SEM #include <semaphore.h>

39233 int sem\_wait(sem\_t \*sem);

39234

39235 **DESCRIPTION**

39236 Refer to *sem\_trywait()*.

39237 **NAME**

39238 semctl — XSI semaphore control operations

39239 **SYNOPSIS**39240 XSI 

```
#include <sys/sem.h>
```

39241 

```
int semctl(int semid, int semnum, int cmd, ...);
```

39242

39243 **DESCRIPTION**

39244 The *semctl()* function operates on XSI semaphores (see the Base Definitions volume of  
 39245 IEEE Std. 1003.1-200x, Section 4.13, Semaphore). It is unspecified whether this function  
 39246 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 39247 page 543).

39248 The *semctl()* function provides a variety of semaphore control operations as specified by *cmd*.  
 39249 The fourth argument is optional and depends upon the operation requested. If required, it is of  
 39250 type **union semun**, which the application shall explicitly declare:

```
39251 union semun {
39252 int val;
39253 struct semid_ds *buf;
39254 unsigned short *array;
39255 }
```

39256 The following semaphore control operations as specified by *cmd* are executed with respect to the  
 39257 semaphore specified by *semid* and *semnum*. The level of permission required for each operation  
 39258 is shown with each command; see Section 2.7 (on page 541). The symbolic names for the values  
 39259 of *cmd* are defined by the `<sys/sem.h>` header:

39260 **GETVAL** Return the value of *semval*; see `<sys/sem.h>`. Requires read permission.

39261 **SETVAL** Set the value of *semval* to *arg.val*, where *arg* is the value of the fourth argument  
 39262 to *semctl()*. When this command is successfully executed, the *semadj* value  
 39263 corresponding to the specified semaphore in all processes is cleared. Requires  
 39264 alter permission; see Section 2.7 (on page 541).

39265 **GETPID** Return the value of *sempid*. Requires read permission.

39266 **GETNCNT** Return the value of *semmcnt*. Requires read permission.

39267 **GETZCNT** Return the value of *semzcnt*. Requires read permission.

39268 The following values of *cmd* operate on each *semval* in the set of semaphores:

39269 **GETALL** Return the value of *semval* for each semaphore in the semaphore set and place  
 39270 into the array pointed to by *arg.array*, where *arg* is the fourth argument to  
 39271 *semctl()*. Requires read permission.

39272 **SETALL** Set the value of *semval* for each semaphore in the semaphore set according to  
 39273 the array pointed to by *arg.array*, where *arg* is the fourth argument to *semctl()*.  
 39274 When this command is successfully executed, the *semadj* values corresponding  
 39275 to each specified semaphore in all processes are cleared. Requires alter  
 39276 permission.

39277 The following values of *cmd* are also available:

39278 **IPC\_STAT** Place the current value of each member of the **semid\_ds** data structure  
 39279 associated with *semid* into the structure pointed to by *arg.buf*, where *arg* is the  
 39280 fourth argument to *semctl()*. The contents of this structure are defined in

|       |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 39281 |                     | <sys/sem.h>. Requires read permission.                                                                                                                                                                                                                                                                                                                                                                                                            |
| 39282 | IPC_SET             | Set the value of the following members of the <b>semid_ds</b> data structure associated with <i>semid</i> to the corresponding value found in the structure pointed to by <i>arg.buf</i> , where <i>arg</i> is the fourth argument to <i>semctl()</i> :                                                                                                                                                                                           |
| 39283 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39284 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39285 |                     | <i>sem_perm.uid</i>                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 39286 |                     | <i>sem_perm.gid</i>                                                                                                                                                                                                                                                                                                                                                                                                                               |
| 39287 |                     | <i>sem_perm.mode</i>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 39288 |                     | The mode bits specified in Section 2.7.1 (on page 541) are copied into the corresponding bits of the <i>sem_perm.mode</i> associated with <i>semid</i> . The stored values of any other bits are unspecified.                                                                                                                                                                                                                                     |
| 39289 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39290 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39291 |                     | This command can only be executed by a process that has an effective user ID equal to either that of a process with appropriate privileges or to the value of <i>sem_perm.cuid</i> or <i>sem_perm.uid</i> in the <b>semid_ds</b> data structure associated with <i>semid</i> .                                                                                                                                                                    |
| 39292 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39293 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39294 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39295 | IPC_RMID            | Remove the semaphore identifier specified by <i>semid</i> from the system and destroy the set of semaphores and <b>semid_ds</b> data structure associated with it. This command can only be executed by a process that has an effective user ID equal to either that of a process with appropriate privileges or to the value of <i>sem_perm.cuid</i> or <i>sem_perm.uid</i> in the <b>semid_ds</b> data structure associated with <i>semid</i> . |
| 39296 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39297 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39298 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39299 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39300 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39301 | <b>RETURN VALUE</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39302 |                     | If successful, the value returned by <i>semctl()</i> depends on <i>cmd</i> as follows:                                                                                                                                                                                                                                                                                                                                                            |
| 39303 | GETVAL              | The value of <i>semval</i> .                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 39304 | GETPID              | The value of <i>sempid</i> .                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 39305 | GETNCNT             | The value of <i>semmcnt</i> .                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 39306 | GETZCNT             | The value of <i>semzcnt</i> .                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 39307 | All others          | 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| 39308 |                     | Otherwise, <i>semctl()</i> shall return $-1$ and set <i>errno</i> to indicate the error.                                                                                                                                                                                                                                                                                                                                                          |
| 39309 | <b>ERRORS</b>       |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39310 |                     | The <i>semctl()</i> function shall fail if:                                                                                                                                                                                                                                                                                                                                                                                                       |
| 39311 | [EACCES]            | Operation permission is denied to the calling process; see Section 2.7 (on page 541).                                                                                                                                                                                                                                                                                                                                                             |
| 39312 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39313 | [EINVAL]            | The value of <i>semid</i> is not a valid semaphore identifier, or the value of <i>semnum</i> is less than 0 or greater than or equal to <i>sem_nsems</i> , or the value of <i>cmd</i> is not a valid command.                                                                                                                                                                                                                                     |
| 39314 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39315 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39316 | [EPERM]             | The argument <i>cmd</i> is equal to IPC_RMID or IPC_SET and the effective user ID of the calling process is not equal to that of a process with appropriate privileges and it is not equal to the value of <i>sem_perm.cuid</i> or <i>sem_perm.uid</i> in the data structure associated with <i>semid</i> .                                                                                                                                       |
| 39317 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39318 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39319 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 39320 | [ERANGE]            | The argument <i>cmd</i> is equal to SETVAL or SETALL and the value to which <i>semval</i> is to be set is greater than the system-imposed maximum.                                                                                                                                                                                                                                                                                                |
| 39321 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

39322 **EXAMPLES**

39323 None.

39324 **APPLICATION USAGE**

39325 The fourth parameter in the SYNOPSIS section is now specified as ". . ." in order to avoid a  
39326 clash with the ISO C standard when referring to the union *semun* (as defined in Issue 3) and for  
39327 backward compatibility.

39328 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
39329 Application developers who need to use IPC should design their applications so that modules  
39330 using the IPC routines described in Section 2.7 (on page 541) can be easily modified to use the  
39331 alternative interfaces.

39332 **RATIONALE**

39333 None.

39334 **FUTURE DIRECTIONS**

39335 None.

39336 **SEE ALSO**

39337 *semget()*, *semop()*, *sem\_close()*, *sem\_destroy()*, *sem\_getvalue()*, *sem\_init()*, *sem\_open()*, *sem\_post()*,  
39338 *sem\_unlink()*, *sem\_wait()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**sys/sem.h**>,  
39339 Section 2.7 (on page 541)

39340 **CHANGE HISTORY**

39341 First released in Issue 2. Derived from Issue 2 of the SVID.

39342 **Issue 4**

39343 The function is no longer marked as OPTIONAL FUNCTIONALITY.

39344 The <**sys/types.h**> and <**sys/ipc.h**> headers are removed from the SYNOPSIS section.

39345 The last argument is now defined by an ellipsis symbol. In previous issues it was defined as a  
39346 union of the various types required by settings of *cmd*. These are now defined individually in  
39347 each description of permitted *cmd* settings. The text of the description of SETALL in the  
39348 DESCRIPTION now refers to the fourth argument instead of *arg.buf*.

39349 In the DESCRIPTION the type of the array is specified in the descriptions of GETALL and  
39350 SETALL.

39351 The [ENOSYS] error is removed from the ERRORS section.

39352 A FUTURE DIRECTIONS section is added warning application developers about migration to  
39353 IEEE 1003.4 interfaces for interprocess communication.

39354 **Issue 4, Version 2**

39355 The fourth argument to *semctl()*, formerly specified in the APPLICATION USAGE section, is  
39356 moved to the DESCRIPTION, and references to its elements are made more precise.

39357 **Issue 5**

39358 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
39359 DIRECTIONS to the APPLICATION USAGE section.

## 39360 NAME

39361 semget — get set of XSI semaphores

## 39362 SYNOPSIS

39363 XSI 

```
#include <sys/sem.h>
```

39364 

```
int semget(key_t key, int nsems, int semflg);
```

39365

## 39366 DESCRIPTION

39367 The *semget()* function operates on XSI semaphores (see the Base Definitions volume of  
 39368 IEEE Std. 1003.1-200x, Section 4.13, Semaphore). It is unspecified whether this function  
 39369 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 39370 page 543).

39371 The *semget()* function shall return the semaphore identifier associated with *key*.

39372 A semaphore identifier with its associated **semid\_ds** data structure and its associated set of  
 39373 *nsems* semaphores (see <sys/sem.h>) is created for *key* if one of the following is true:

- 39374 • The argument *key* is equal to `IPC_PRIVATE`.
- 39375 • The argument *key* does not already have a semaphore identifier associated with it and (*semflg*  
 39376 & `IPC_CREAT`) is non-zero.

39377 Upon creation, the **semid\_ds** data structure associated with the new semaphore identifier is  
 39378 initialized as follows:

- 39379 • In the operation permissions structure *sem\_perm.cuid*, *sem\_perm.uid*, *sem\_perm.cgid*, and  
 39380 *sem\_perm.gid* are set equal to the effective user ID and effective group ID, respectively, of the  
 39381 calling process.
- 39382 • The low-order 9 bits of *sem\_perm.mode* are set equal to the low-order 9 bits of *semflg*.
- 39383 • The variable *sem\_nsems* is set equal to the value of *nsems*.
- 39384 • The variable *sem\_otime* is set equal to 0 and *sem\_ctime* is set equal to the current time.
- 39385 • The data structure associated with each semaphore in the set is not initialized. The *semctl()*  
 39386 function with the command `SETVAL` or `SETALL` can be used to initialize each semaphore.

## 39387 RETURN VALUE

39388 Upon successful completion, *semget()* shall return a non-negative integer, namely a semaphore  
 39389 identifier; otherwise, it shall return `-1` and set *errno* to indicate the error.

## 39390 ERRORS

39391 The *semget()* function shall fail if:

- |                                  |          |                                                                                                                                                                                                                                                                                               |
|----------------------------------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 39392<br>39393<br>39394          | [EACCES] | A semaphore identifier exists for <i>key</i> , but operation permission as specified by the low-order 9 bits of <i>semflg</i> would not be granted; see Section 2.7 (on page 541).                                                                                                            |
| 39395<br>39396                   | [EEXIST] | A semaphore identifier exists for the argument <i>key</i> but (( <i>semflg</i> & <code>IPC_CREAT</code> ) && ( <i>semflg</i> & <code>IPC_EXCL</code> )) is non-zero.                                                                                                                          |
| 39397<br>39398<br>39399<br>39400 | [EINVAL] | The value of <i>nsems</i> is either less than or equal to 0 or greater than the system-imposed limit, or a semaphore identifier exists for the argument <i>key</i> , but the number of semaphores in the set associated with it is less than <i>nsems</i> and <i>nsems</i> is not equal to 0. |
| 39401<br>39402                   | [ENOENT] | A semaphore identifier does not exist for the argument <i>key</i> and ( <i>semflg</i> & <code>IPC_CREAT</code> ) is equal to 0.                                                                                                                                                               |

39403 [ENOSPC] A semaphore identifier is to be created but the system-imposed limit on the  
 39404 maximum number of allowed semaphores system-wide would be exceeded.

### 39405 EXAMPLES

#### 39406 Creating a Semaphore Identifier

39407 The following example gets a unique semaphore key using the *ftok()* function, then gets a  
 39408 semaphore ID associated with that key using the *semget()* function (the first call also tests to  
 39409 make sure the semaphore exists). If the semaphore does not exist, the program creates it, as  
 39410 shown by the second call to *semget()*. In creating the semaphore for the queueing process, the  
 39411 program attempts to create one semaphore with read/write permission for all. It also uses the  
 39412 *IPC\_EXCL* flag, which forces *semget()* to fail if the semaphore already exists.

39413 After creating the semaphore, the program uses a call to *semop()* to initialize it to the values in  
 39414 the *sbuf* array. The number of processes that can execute concurrently without queuing is  
 39415 initially set to 2. The final call to *semget()* creates a semaphore identifier that can be used later in  
 39416 the program.

```

39417 #include <sys/types.h>
39418 #include <stdio.h>
39419 #include <sys/ipc.h>
39420 #include <sys/sem.h>
39421 #include <sys/stat.h>
39422 #include <errno.h>
39423 #include <unistd.h>
39424 #include <stdlib.h>
39425 #include <pwd.h>
39426 #include <fcntl.h>
39427 #include <limits.h>
39428 ...
39429 key_t semkey;
39430 int semid, pfd, fv;
39431 struct sembuf sbuf;
39432 char *lgn;
39433 char filename[PATH_MAX+1];
39434 struct stat outstat;
39435 struct passwd *pw;
39436 ...
39437 /* Get unique key for semaphore. */
39438 if ((semkey = ftok("/tmp", 'a')) == (key_t) -1) {
39439 perror("IPC error: ftok"); exit(1);
39440 }
39441 /* Get semaphore ID associated with this key. */
39442 if ((semid = semget(semkey, 0, 0)) == -1) {
39443 /* Semaphore does not exist - Create. */
39444 if ((semid = semget(semkey, 1, IPC_CREAT | IPC_EXCL | S_IRUSR |
39445 S_IWUSR | S_IRGRP | S_IWGRP | S_IROTH | S_IWOTH)) != -1)
39446 {
39447 /* Initialize the semaphore. */
39448 sbuf.sem_num = 0;
39449 sbuf.sem_op = 2; /* This is the number of runs without queuing. */
39450 sbuf.sem_flg = 0;

```

```

39451 if (semop(semid, &sbuf, 1) == -1) {
39452 perror("IPC error: semop"); exit(1);
39453 }
39454 }
39455 else if (errno == EEXIST) {
39456 if ((semid = semget(semkey, 0, 0)) == -1) {
39457 perror("IPC error 1: semget"); exit(1);
39458 }
39459 }
39460 else {
39461 perror("IPC error 2: semget"); exit(1);
39462 }
39463 }
39464 ...

```

#### 39465 APPLICATION USAGE

39466 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
 39467 Application developers who need to use IPC should design their applications so that modules  
 39468 using the IPC routines described in Section 2.7 (on page 541) can be easily modified to use the  
 39469 alternative interfaces.

#### 39470 RATIONALE

39471 None.

#### 39472 FUTURE DIRECTIONS

39473 None.

#### 39474 SEE ALSO

39475 *semctl()*, *semop()*, *sem\_close()*, *sem\_destroy()*, *sem\_getvalue()*, *sem\_init()*, *sem\_open()*, *sem\_post()*,  
 39476 *sem\_unlink()*, *sem\_wait()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<sys/sem.h>`,  
 39477 Section 2.7 (on page 541).

#### 39478 CHANGE HISTORY

39479 First released in Issue 2. Derived from Issue 2 of the SVID.

#### 39480 Issue 4

39481 The function is no longer marked as OPTIONAL FUNCTIONALITY.

39482 The `<sys/types.h>` and `<sys/ipc.h>` headers are removed from the SYNOPSIS section.

39483 The [ENOSYS] error is removed from the ERRORS section.

39484 A FUTURE DIRECTIONS section is added warning application developers about migration to  
 39485 IEEE 1003.4 interfaces for interprocess communication.

#### 39486 Issue 5

39487 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
 39488 DIRECTIONS to a new APPLICATION USAGE section.

39489 **NAME**

39490 semop — XSI semaphore operations

39491 **SYNOPSIS**39492 XSI 

```
#include <sys/sem.h>
```

39493 

```
int semop(int semid, struct sembuf *sops, size_t nsops);
```

39494

39495 **DESCRIPTION**

39496 The *semop()* function operates on XSI semaphores (see the Base Definitions volume of  
 39497 IEEE Std. 1003.1-200x, Section 4.13, Semaphore). It is unspecified whether this function  
 39498 interoperates with the realtime interprocess communication facilities defined in Section 2.8 (on  
 39499 page 543).

39500 The *semop()* function is used to perform atomically a user-defined array of semaphore  
 39501 operations on the set of semaphores associated with the semaphore identifier specified by the  
 39502 argument *semid*.

39503 The argument *sops* is a pointer to a user-defined array of semaphore operation structures. The  
 39504 implementation shall not modify elements of this array unless the application uses  
 39505 implementation-defined extensions.

39506 The argument *nsops* is the number of such structures in the array.

39507 Each structure, **sembuf**, includes the following members:

39508

39509

| Member Type | Member Name    | Description          |
|-------------|----------------|----------------------|
| short       | <i>sem_num</i> | Semaphore number.    |
| short       | <i>sem_op</i>  | Semaphore operation. |
| short       | <i>sem_flg</i> | Operation flags.     |

39510

39511

39512

39513 Each semaphore operation specified by *sem\_op* is performed on the corresponding semaphore  
 39514 specified by *semid* and *sem\_num*.

39515 The variable *sem\_op* specifies one of three semaphore operations:

39516 1. If *sem\_op* is a negative integer and the calling process has alter permission, one of the  
 39517 following shall occur:

39518 • If *semval* (see **<sys/sem.h>**) is greater than or equal to the absolute value of *sem\_op*, the  
 39519 absolute value of *sem\_op* is subtracted from *semval*. Also, if (*sem\_flg* & SEM\_UNDO) is  
 39520 non-zero, the absolute value of *sem\_op* is added to the calling process' *semadj* value for  
 39521 the specified semaphore.

39522 • If *semval* is less than the absolute value of *sem\_op* and (*sem\_flg* & IPC\_NOWAIT) is non-  
 39523 zero, *semop()* shall return immediately.

39524 • If *semval* is less than the absolute value of *sem\_op* and (*sem\_flg* & IPC\_NOWAIT) is 0,  
 39525 *semop()* shall increment the *semncnt* associated with the specified semaphore and  
 39526 suspend execution of the calling thread until one of the following conditions occurs:

39527 — The value of *semval* becomes greater than or equal to the absolute value of *sem\_op*.  
 39528 When this occurs, the value of *semncnt* associated with the specified semaphore is  
 39529 decremented, the absolute value of *sem\_op* is subtracted from *semval* and, if (*sem\_flg*  
 39530 & SEM\_UNDO) is non-zero, the absolute value of *sem\_op* is added to the calling  
 39531 process' *semadj* value for the specified semaphore.

- 39532 — The *semid* for which the calling thread is awaiting action is removed from the  
 39533 system. When this occurs, *errno* shall be set equal to [EIDRM] and  $-1$  shall be  
 39534 returned.
- 39535 — The calling thread receives a signal that is to be caught. When this occurs, the value  
 39536 of *semncnt* associated with the specified semaphore is decremented, and the calling  
 39537 thread resumes execution in the manner prescribed in *sigaction()*.
- 39538 2. If *sem\_op* is a positive integer and the calling process has alter permission, the value of  
 39539 *sem\_op* is added to *semval* and, if (*sem\_flg* & SEM\_UNDO) is non-zero, the value of *sem\_op* is  
 39540 subtracted from the calling process' *semadj* value for the specified semaphore.
- 39541 3. If *sem\_op* is 0 and the calling process has read permission, one of the following shall occur:
- 39542 • If *semval* is 0, *semop()* shall return immediately.
- 39543 • If *semval* is non-zero and (*sem\_flg* & IPC\_NOWAIT) is non-zero, *semop()* shall return  
 39544 immediately.
- 39545 • If *semval* is non-zero and (*sem\_flg* & IPC\_NOWAIT) is 0, *semop()* will increment the  
 39546 *semzcnt* associated with the specified semaphore and suspends execution of the calling  
 39547 thread until one of the following occurs:
- 39548 — The value of *semval* becomes 0, at which time the value of *semzcnt* associated with  
 39549 the specified semaphore is decremented.
- 39550 — The *semid* for which the calling thread is awaiting action is removed from the  
 39551 system. When this occurs, *errno* shall be set equal to [EIDRM] and  $-1$  shall be  
 39552 returned.
- 39553 — The calling thread receives a signal that is to be caught. When this occurs, the value  
 39554 of *semzcnt* associated with the specified semaphore is decremented, and the calling  
 39555 thread resumes execution in the manner prescribed in *sigaction()*.
- 39556 Upon successful completion, the value of *sempid* for each semaphore specified in the array  
 39557 pointed to by *sops* shall be set equal to the process ID of the calling process.

39558 **RETURN VALUE**

39559 Upon successful completion, *semop()* shall return 0; otherwise, it shall return  $-1$  and set *errno* to  
 39560 indicate the error.

39561 **ERRORS**

39562 The *semop()* function shall fail if:

|       |          |                                                                                      |  |
|-------|----------|--------------------------------------------------------------------------------------|--|
| 39563 | [E2BIG]  | The value of <i>nsops</i> is greater than the system-imposed maximum.                |  |
| 39564 | [EACCES] | Operation permission is denied to the calling process; see Section 2.7 (on page      |  |
| 39565 |          | 541).                                                                                |  |
| 39566 | [EAGAIN] | The operation would result in suspension of the calling process but ( <i>sem_flg</i> |  |
| 39567 |          | & IPC_NOWAIT) is non-zero.                                                           |  |
| 39568 | [EFBIG]  | The value of <i>sem_num</i> is less than 0 or greater than or equal to the number of |  |
| 39569 |          | semaphores in the set associated with <i>semid</i> .                                 |  |
| 39570 | [EIDRM]  | The semaphore identifier <i>semid</i> is removed from the system.                    |  |
| 39571 | [EINTR]  | The <i>semop()</i> function was interrupted by a signal.                             |  |
| 39572 | [EINVAL] | The value of <i>semid</i> is not a valid semaphore identifier, or the number of      |  |
| 39573 |          | individual semaphores for which the calling process requests a SEM_UNDO              |  |
| 39574 |          | would exceed the system-imposed limit.                                               |  |

39575 [ENOSPC] The limit on the number of individual processes requesting a SEM\_UNDO |  
 39576 would be exceeded.

39577 [ERANGE] An operation would cause a *semval* to overflow the system-imposed limit, or |  
 39578 an operation would cause a *semadj* value to overflow the system-imposed |  
 39579 limit.

### 39580 EXAMPLES

#### 39581 Setting Values in Semaphores

39582 The following example sets the values of the two semaphores associated with the *semid*  
 39583 identifier to the values contained in the *sb* array.

```
39584 #include <sys/sem.h>
39585 ...
39586 int semid;
39587 struct sembuf sb[2];
39588 int nsops = 2;
39589 int result;

39590 /* Adjust value of semaphore in the semaphore array semid. */
39591 sb[0].sem_num = 0;
39592 sb[0].sem_op = -1;
39593 sb[0].sem_flg = SEM_UNDO | IPC_NOWAIT;
39594 sb[1].sem_num = 1;
39595 sb[1].sem_op = 1;
39596 sb[1].sem_flg = 0;

39597 result = semop(semid, sb, nsops);
```

#### 39598 Creating a Semaphore Identifier

39599 The following example gets a unique semaphore key using the *ftok()* function, then gets a  
 39600 semaphore ID associated with that key using the *semget()* function (the first call also tests to  
 39601 make sure the semaphore exists). If the semaphore does not exist, the program creates it, as  
 39602 shown by the second call to *semget()*. In creating the semaphore for the queueing process, the  
 39603 program attempts to create one semaphore with read/write permission for all. It also uses the  
 39604 IPC\_EXCL flag, which forces *semget()* to fail if the semaphore already exists.

39605 After creating the semaphore, the program uses a call to *semop()* to initialize it to the values in  
 39606 the *sbuf* array. The number of processes that can execute concurrently without queuing is  
 39607 initially set to 2. The final call to *semget()* creates a semaphore identifier that can be used later in  
 39608 the program.

39609 The final call to *semop()* acquires the semaphore and waits until it is free; the SEM\_UNDO  
 39610 option releases the semaphore when the process exits, waiting until there are less than two  
 39611 processes running concurrently.

```
39612 #include <sys/types.h>
39613 #include <stdio.h>
39614 #include <sys/ipc.h>
39615 #include <sys/sem.h>
39616 #include <sys/stat.h>
39617 #include <errno.h>
39618 #include <unistd.h>
39619 #include <stdlib.h>
```

```

39620 #include <pwd.h>
39621 #include <fcntl.h>
39622 #include <limits.h>
39623 ...
39624 key_t semkey;
39625 int semid, pfd, fv;
39626 struct sembuf sbuf;
39627 char *lgn;
39628 char filename[PATH_MAX+1];
39629 struct stat outstat;
39630 struct passwd *pw;
39631 ...
39632 /* Get unique key for semaphore. */
39633 if ((semkey = ftok("/tmp", 'a')) == (key_t) -1) {
39634 perror("IPC error: ftok"); exit(1);
39635 }
39636 /* Get semaphore ID associated with this key. */
39637 if ((semid = semget(semkey, 0, 0)) == -1) {
39638 /* Semaphore does not exist - Create. */
39639 if ((semid = semget(semkey, 1, IPC_CREAT | IPC_EXCL | S_IRUSR |
39640 S_IWUSR | S_IRGRP | S_IWGRP | S_IROTH | S_IWOTH)) != -1)
39641 {
39642 /* Initialize the semaphore. */
39643 sbuf.sem_num = 0;
39644 sbuf.sem_op = 2; /* This is the number of runs without queuing. */
39645 sbuf.sem_flg = 0;
39646 if (semop(semid, &sbuf, 1) == -1) {
39647 perror("IPC error: semop"); exit(1);
39648 }
39649 }
39650 else if (errno == EEXIST) {
39651 if ((semid = semget(semkey, 0, 0)) == -1) {
39652 perror("IPC error 1: semget"); exit(1);
39653 }
39654 }
39655 else {
39656 perror("IPC error 2: semget"); exit(1);
39657 }
39658 }
39659 ...
39660 sbuf.sem_num = 0;
39661 sbuf.sem_op = -1;
39662 sbuf.sem_flg = SEM_UNDO;
39663 if (semop(semid, &sbuf, 1) == -1) {
39664 perror("IPC Error: semop"); exit(1);
39665 }

```

#### 39666 APPLICATION USAGE

39667 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
39668 Application developers who need to use IPC should design their applications so that modules  
39669 using the IPC routines described in Section 2.7 (on page 541) can be easily modified to use the  
39670 alternative interfaces.

39671 **RATIONALE**

39672 None.

39673 **FUTURE DIRECTIONS**

39674 None.

39675 **SEE ALSO**

39676 *exec*, *exit()*, *fork()*, *semctl()*, *semget()*, *sem\_close()*, *sem\_destroy()*, *sem\_getvalue()*, *sem\_init()*,  
39677 *sem\_open()*, *sem\_post()*, *sem\_unlink()*, *sem\_wait()*, the Base Definitions volume of  
39678 IEEE Std. 1003.1-200x, <sys/ipc.h>, <sys/sem.h>, <sys/types.h>, Section 2.7 (on page 541)

39679 **CHANGE HISTORY**

39680 First released in Issue 2. Derived from Issue 2 of the SVID.

39681 **Issue 4**

39682 The function is no longer marked as OPTIONAL FUNCTIONALITY.

39683 The &lt;sys/types.h&gt; and &lt;sys/ipc.h&gt; headers are removed from the SYNOPSIS section.

39684 The type of *nsops* is changed to **size\_t**.

39685 The DESCRIPTION is updated to indicate that an implementation does not modify the elements  
39686 of *sops* unless the application uses implementation-defined extensions.

39687 The [ENOSYS] error is removed from the ERRORS section.

39688 A FUTURE DIRECTIONS section is added warning application developers about migration to  
39689 IEEE 1003.4 interfaces for interprocess communication.

39690 **Issue 5**

39691 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
39692 DIRECTIONS to a new APPLICATION USAGE section.

## 39693 NAME

39694 send — send a message on a socket

## 39695 SYNOPSIS

39696 #include &lt;sys/socket.h&gt;

39697 ssize\_t send(int *socket*, const void \**buffer*, size\_t *length*, int *flags*);

## 39698 DESCRIPTION

39699 The *send()* functions takes the following arguments:

|       |               |                                                                                                                                                      |
|-------|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| 39700 | <i>socket</i> | Specifies the socket file descriptor.                                                                                                                |
| 39701 | <i>buffer</i> | Points to the buffer containing the message to send.                                                                                                 |
| 39702 | <i>length</i> | Specifies the length of the message in bytes.                                                                                                        |
| 39703 | <i>flags</i>  | Specifies the type of message transmission. Values of this argument are formed by logically OR'ing zero or more of the following flags:              |
| 39704 |               |                                                                                                                                                      |
| 39705 | MSG_EOR       | Terminates a record (if supported by the protocol).                                                                                                  |
| 39706 | MSG_OOB       | Sends out-of-band data on sockets that support out-of-band communications. The significance and semantics of out-of-band data are protocol-specific. |
| 39707 |               |                                                                                                                                                      |
| 39708 |               |                                                                                                                                                      |

39709 The *send()* function initiates transmission of a message from the specified socket to its peer. The  
 39710 *send()* function sends a message only when the socket is connected (including when the peer of a  
 39711 connectionless socket has been set via *connect()*).

39712 The length of the message to be sent is specified by the *length* argument. If the message is too  
 39713 long to pass through the underlying protocol, *send()* shall fail and no data shall be transmitted.

39714 Successful completion of a call to *send()* does not guarantee delivery of the message. A return  
 39715 value of  $-1$  indicates only locally-detected errors.

39716 If space is not available at the sending socket to hold the message to be transmitted, and the  
 39717 socket file descriptor does not have O\_NONBLOCK set, *send()* shall block until space is  
 39718 available. If space is not available at the sending socket to hold the message to be transmitted,  
 39719 and the socket file descriptor does have O\_NONBLOCK set, *send()* shall fail. The *select()* and  
 39720 *poll()* functions can be used to determine when it is possible to send more data.

39721 The socket in use may require the process to have appropriate privileges to use the *send()*  
 39722 function.

## 39723 RETURN VALUE

39724 Upon successful completion, *send()* shall return the number of bytes sent. Otherwise,  $-1$  shall be  
 39725 returned and *errno* set to indicate the error.

## 39726 ERRORS

39727 The *send()* function shall fail if:

|       |                           |                                                                                            |
|-------|---------------------------|--------------------------------------------------------------------------------------------|
| 39728 | [EAGAIN] or [EWOULDBLOCK] |                                                                                            |
| 39729 |                           | The socket's file descriptor is marked O_NONBLOCK and the requested operation would block. |
| 39730 |                           |                                                                                            |
| 39731 | [EBADF]                   | The <i>socket</i> argument is not a valid file descriptor.                                 |
| 39732 | [ECONNRESET]              | A connection was forcibly closed by a peer.                                                |
| 39733 | [EDESTADDRREQ]            |                                                                                            |
| 39734 |                           | The socket is not connection-mode and no peer address is set.                              |

- 39735 [EINTR] A signal interrupted *send()* before any data was transmitted.
- 39736 [EMSGSIZE] The message is too large be sent all at once, as the socket requires.
- 39737 [ENOTCONN] The socket is not connected or otherwise has not had the peer pre-specified.
- 39738 [ENOTSOCK] The *socket* argument does not refer to a socket.
- 39739 [EOPNOTSUPP] The *socket* argument is associated with a socket that does not support one or  
39740 more of the values set in *flags*.
- 39741 [EPIPE] The socket is shut down for writing, or the socket is connection-mode and is  
39742 no longer connected. In the latter case, and if the socket is of type  
39743 SOCK\_STREAM, the SIGPIPE signal is generated to the calling thread.
- 39744 The *send()* function may fail if:
- 39745 [EACCES] The calling process does not have the appropriate privileges.
- 39746 [EIO] An I/O error occurred while reading from or writing to the file system.
- 39747 [ENETDOWN] The local network interface used to reach the destination is down.
- 39748 [ENETUNREACH]  
39749 No route to the network is present.
- 39750 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 39751 **EXAMPLES**
- 39752 None.
- 39753 **APPLICATION USAGE**
- 39754 The *send()* function is identical to *sendto()* with a null pointer *dest\_len* argument, and to *write()* if  
39755 no flags are used.
- 39756 **RATIONALE**
- 39757 None.
- 39758 **FUTURE DIRECTIONS**
- 39759 None.
- 39760 **SEE ALSO**
- 39761 *connect()*, *getsockopt()*, *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *sendmsg()*, *sendto()*,  
39762 *setsockopt()*, *shutdown()*, *socket()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
39763 <sys/socket.h>
- 39764 **CHANGE HISTORY**
- 39765 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

## 39766 NAME

39767 sendmsg — send a message on a socket using a message structure

## 39768 SYNOPSIS

39769 #include <sys/socket.h>

39770 ssize\_t sendmsg(int *socket*, const struct msghdr \**message*, int *flags*);

## 39771 DESCRIPTION

39772 The *sendmsg()* function sends a message through a connection-mode or connectionless-mode  
 39773 socket. If the socket is connectionless-mode, the message shall be sent to the address specified by  
 39774 **msghdr**. If the socket is connection-mode, the destination address in **msghdr** is ignored.

39775 The *sendmsg()* function takes the following arguments:

|       |                |                                                                                                                                                                                                                                                 |
|-------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 39776 | <i>socket</i>  | Specifies the socket file descriptor.                                                                                                                                                                                                           |
| 39777 | <i>message</i> | Points to a <b>msghdr</b> structure, containing both the destination address and the buffers for the outgoing message. The length and format of the address depend on the address family of the socket. The <i>msg_flags</i> member is ignored. |
| 39780 | <i>flags</i>   | Specifies the type of message transmission. The application may specify 0 or the following flag:                                                                                                                                                |
| 39782 | MSG_EOR        | Terminates a record (if supported by the protocol).                                                                                                                                                                                             |
| 39783 | MSG_OOB        | Sends out-of-band data on sockets that support out-of-band data. The significance and semantics of out-of-band data are protocol-specific.                                                                                                      |

39786 The *msg\_iov* and *msg\_iovlen* fields of *message* specify zero or more buffers containing the data to  
 39787 be sent. *msg\_iov* points to an array of **iovec** structures; *msg\_iovlen* shall be set to the dimension of  
 39788 this array. In each **iovec** structure, the *iov\_base* field specifies a storage area and the *iov\_len* field  
 39789 gives its size in bytes. Some of these sizes can be zero. The data from each storage area indicated  
 39790 by *msg\_iov* is sent in turn.

39791 Successful completion of a call to *sendmsg()* does not guarantee delivery of the message. A  
 39792 return value of  $-1$  indicates only locally-detected errors.

39793 If space is not available at the sending socket to hold the message to be transmitted and the  
 39794 socket file descriptor does not have O\_NONBLOCK set, *sendmsg()* function blocks until space is  
 39795 available. If space is not available at the sending socket to hold the message to be transmitted  
 39796 and the socket file descriptor does have O\_NONBLOCK set, *sendmsg()* function shall fail.

39797 If the socket protocol supports broadcast and the specified address is a broadcast address for the  
 39798 socket protocol, *sendmsg()* shall fail if the SO\_BROADCAST option is not set for the socket.

39799 The socket in use may require the process to have appropriate privileges to use the *sendmsg()*  
 39800 function.

## 39801 RETURN VALUE

39802 Upon successful completion, *sendmsg()* shall return the number of bytes sent. Otherwise,  $-1$   
 39803 shall be returned and *errno* set to indicate the error.

## 39804 ERRORS

39805 The *sendmsg()* function shall fail if:

39806 [EAGAIN] or [EWOULDBLOCK]

39807 The socket's file descriptor is marked O\_NONBLOCK and the requested  
 39808 operation would block.

|       |                |                                                                                                                                                                                     |
|-------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 39809 | [EAFNOSUPPORT] |                                                                                                                                                                                     |
| 39810 |                | Addresses in the specified address family cannot be used with this socket.                                                                                                          |
| 39811 | [EBADF]        | The <i>socket</i> argument is not a valid file descriptor.                                                                                                                          |
| 39812 | [ECONNRESET]   | A connection was forcibly closed by a peer.                                                                                                                                         |
| 39813 | [EINTR]        | A signal interrupted <i>sendmsg()</i> before any data was transmitted.                                                                                                              |
| 39814 | [EINVAL]       | The sum of the <i>iov_len</i> values overflows an <b>ssize_t</b> .                                                                                                                  |
| 39815 | [EMSGSIZE]     | The message is too large to be sent all at once (as the socket requires), or the <i>msg_iovlen</i> member of the <b>msg_hdr</b> structure pointed to by <i>message</i> is less than |
| 39816 |                | or equal to 0 or is greater than {IOV_MAX}.                                                                                                                                         |
| 39817 |                |                                                                                                                                                                                     |
| 39818 | [ENOTCONN]     | The socket is connection-mode but is not connected.                                                                                                                                 |
| 39819 | [ENOTSOCK]     | The <i>socket</i> argument does not refer a socket.                                                                                                                                 |
| 39820 | [EOPNOTSUPP]   | The <i>socket</i> argument is associated with a socket that does not support one or                                                                                                 |
| 39821 |                | more of the values set in <i>flags</i> .                                                                                                                                            |
| 39822 | [EPIPE]        | The socket is shut down for writing, or the socket is connection-mode and is                                                                                                        |
| 39823 |                | no longer connected. In the latter case, and if the socket is of type                                                                                                               |
| 39824 |                | SOCK_STREAM, the SIGPIPE signal is generated to the calling thread.                                                                                                                 |
| 39825 |                | If the address family of the socket is AF_UNIX, then <i>sendmsg()</i> shall fail if:                                                                                                |
| 39826 | [EIO]          | An I/O error occurred while reading from or writing to the file system.                                                                                                             |
| 39827 | [ELOOP]        | A loop exists in symbolic links encountered during resolution of the path                                                                                                           |
| 39828 |                | name in the socket address.                                                                                                                                                         |
| 39829 | [ENAMETOOLONG] |                                                                                                                                                                                     |
| 39830 |                | A component of a path name exceeded {NAME_MAX} characters, or an entire                                                                                                             |
| 39831 |                | path name exceeded {PATH_MAX} characters.                                                                                                                                           |
| 39832 | [ENOENT]       | A component of the path name does not name an existing file or the path                                                                                                             |
| 39833 |                | name is an empty string.                                                                                                                                                            |
| 39834 | [ENOTDIR]      | A component of the path prefix of the path name in the socket address is not a                                                                                                      |
| 39835 |                | directory.                                                                                                                                                                          |
| 39836 |                | The <i>sendmsg()</i> function may fail if:                                                                                                                                          |
| 39837 | [EACCES]       | Search permission is denied for a component of the path prefix; or write                                                                                                            |
| 39838 |                | access to the named socket is denied.                                                                                                                                               |
| 39839 | [EDESTADDRREQ] |                                                                                                                                                                                     |
| 39840 |                | The socket is not connection-mode and does not have its peer address set, and                                                                                                       |
| 39841 |                | no destination address was specified.                                                                                                                                               |
| 39842 | [EHOSTUNREACH] |                                                                                                                                                                                     |
| 39843 |                | The destination host cannot be reached (probably because the host is down or                                                                                                        |
| 39844 |                | a remote router cannot reach it).                                                                                                                                                   |
| 39845 | [EIO]          | An I/O error occurred while reading from or writing to the file system.                                                                                                             |
| 39846 | [EISCONN]      | A destination address was specified and the socket is already connected.                                                                                                            |
| 39847 | [ENETDOWN]     | The local network interface used to reach the destination is down.                                                                                                                  |
| 39848 | [ENETUNREACH]  |                                                                                                                                                                                     |
| 39849 |                | No route to the network is present.                                                                                                                                                 |

- 39850 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 39851 [ENOMEM] Insufficient memory was available to fulfill the request.
- 39852 If the address family of the socket is AF\_UNIX, then *sendmsg()* may fail if:
- 39853 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
39854 resolution of the path name in the socket address.
- 39855 [ENAMETOOLONG]  
39856 Path name resolution of a symbolic link produced an intermediate result  
39857 whose length exceeds {PATH\_MAX}.
- 39858 **EXAMPLES**  
39859 Done.
- 39860 **APPLICATION USAGE**  
39861 The *select()* and *poll()* functions can be used to determine when it is possible to send more data.
- 39862 **RATIONALE**  
39863 None.
- 39864 **FUTURE DIRECTIONS**  
39865 None.
- 39866 **SEE ALSO**  
39867 *getsockopt()*, *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendto()*, *setsockopt()*,  
39868 *shutdown()*, *socket()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**sys/socket.h**>
- 39869 **CHANGE HISTORY**  
39870 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.
- 39871 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
39872 [ELOOP] error condition is added.

39873 **NAME**39874 `sendto` — send a message on a socket39875 **SYNOPSIS**39876 `#include <sys/socket.h>`39877 `ssize_t sendto(int socket, const void *message, size_t length,`39878 `int flags, const struct sockaddr *dest_addr,`39879 `socklen_t dest_len);`39880 **DESCRIPTION**

39881 The `sendto()` function sends a message through a connection-mode or connectionless-mode  
 39882 socket. If the socket is connectionless-mode, the message shall be sent to the address specified by  
 39883 `dest_addr`. If the socket is connection-mode, `dest_addr` is ignored.

39884 The `sendto()` function takes the following arguments:

|       |                        |                                                                                                                                                                     |
|-------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 39885 | <code>socket</code>    | Specifies the socket file descriptor.                                                                                                                               |
| 39886 | <code>message</code>   | Points to a buffer containing the message to be sent.                                                                                                               |
| 39887 | <code>length</code>    | Specifies the size of the message in bytes.                                                                                                                         |
| 39888 | <code>flags</code>     | Specifies the type of message transmission. Values of this argument are<br>39889 formed by logically OR'ing zero or more of the following flags:                    |
| 39890 | <code>MSG_EOR</code>   | Terminates a record (if supported by the protocol).                                                                                                                 |
| 39891 | <code>MSG_OOB</code>   | Sends out-of-band data on sockets that support out-of-band<br>39892 data. The significance and semantics of out-of-band data are<br>39893 protocol-specific.        |
| 39894 | <code>dest_addr</code> | Points to a <b>sockaddr</b> structure containing the destination address. The length<br>39895 and format of the address depend on the address family of the socket. |
| 39896 | <code>dest_len</code>  | Specifies the length of the <b>sockaddr</b> structure pointed to by the <code>dest_addr</code><br>39897 argument.                                                   |

39898 If the socket protocol supports broadcast and the specified address is a broadcast address for the  
 39899 socket protocol, `sendto()` shall fail if the `SO_BROADCAST` option is not set for the socket.

39900 The `dest_addr` argument specifies the address of the target. The `length` argument specifies the  
 39901 length of the message.

39902 Successful completion of a call to `sendto()` does not guarantee delivery of the message. A return  
 39903 value of `-1` indicates only locally-detected errors.

39904 If space is not available at the sending socket to hold the message to be transmitted and the  
 39905 socket file descriptor does not have `O_NONBLOCK` set, `sendto()` blocks until space is available.  
 39906 If space is not available at the sending socket to hold the message to be transmitted and the  
 39907 socket file descriptor does have `O_NONBLOCK` set, `sendto()` shall fail.

39908 The socket in use may require the process to have appropriate privileges to use the `sendto()`  
 39909 function.

39910 **RETURN VALUE**

39911 Upon successful completion, `sendto()` shall return the number of bytes sent. Otherwise, `-1` shall  
 39912 be returned and `errno` set to indicate the error.

39913 **ERRORS**

- 39914 The *sendto()* function shall fail if:
- 39915 [EAFNOSUPPORT]  
39916 Addresses in the specified address family cannot be used with this socket.
- 39917 [EAGAIN] or [EWOULDBLOCK]  
39918 The socket's file descriptor is marked O\_NONBLOCK and the requested  
39919 operation would block.
- 39920 [EBADF] The *socket* argument is not a valid file descriptor.
- 39921 [ECONNRESET] A connection was forcibly closed by a peer.
- 39922 [EINTR] A signal interrupted *sendto()* before any data was transmitted.
- 39923 [EMSGSIZE] The message is too large to be sent all at once, as the socket requires.
- 39924 [ENOTCONN] The socket is connection-mode but is not connected.
- 39925 [ENOTSOCK] The *socket* argument does not refer to a socket.
- 39926 [EOPNOTSUPP] The *socket* argument is associated with a socket that does not support one or  
39927 more of the values set in *flags*.
- 39928 [EPIPE] The socket is shut down for writing, or the socket is connection-mode and is  
39929 no longer connected. In the latter case, and if the socket is of type  
39930 SOCK\_STREAM, the SIGPIPE signal is generated to the calling thread.
- 39931 If the address family of the socket is AF\_UNIX, then *sendto()* shall fail if:
- 39932 [EIO] An I/O error occurred while reading from or writing to the file system.
- 39933 [ELOOP] A loop exists in symbolic links encountered during resolution of the path  
39934 name in the socket address.
- 39935 [ENAMETOOLONG]  
39936 A component of a path name exceeded {NAME\_MAX} characters, or an entire  
39937 path name exceeded {PATH\_MAX} characters.
- 39938 [ENOENT] A component of the path name does not name an existing file or the path  
39939 name is an empty string.
- 39940 [ENOTDIR] A component of the path prefix of the path name in the socket address is not a  
39941 directory.
- 39942 The *sendto()* function may fail if:
- 39943 [EACCES] Search permission is denied for a component of the path prefix; or write  
39944 access to the named socket is denied.
- 39945 [EDESTADDRREQ]  
39946 The socket is not connection-mode and does not have its peer address set, and  
39947 no destination address was specified.
- 39948 [EHOSTUNREACH]  
39949 The destination host cannot be reached (probably because the host is down or  
39950 a remote router cannot reach it).
- 39951 [EINVAL] The *dest\_len* argument is not a valid length for the address family.
- 39952 [EIO] An I/O error occurred while reading from or writing to the file system.

- 39953 [EISCONN] A destination address was specified and the socket is already connected. This  
39954 error may or may not be returned for connection mode sockets.
- 39955 [ENETDOWN] The local network interface used to reach the destination is down. |
- 39956 [ENETUNREACH]  
39957 No route to the network is present.
- 39958 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 39959 [ENOMEM] Insufficient memory was available to fulfill the request. |
- 39960 If the address family of the socket is AF\_UNIX, then *sendto()* may fail if:
- 39961 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during |  
39962 resolution of the path name in the socket address. |
- 39963 [ENAMETOOLONG]  
39964 Path name resolution of a symbolic link produced an intermediate result  
39965 whose length exceeds {PATH\_MAX}.
- 39966 **EXAMPLES**  
39967 None.
- 39968 **APPLICATION USAGE**  
39969 The *select()* and *poll()* functions can be used to determine when it is possible to send more data.
- 39970 **RATIONALE**  
39971 None.
- 39972 **FUTURE DIRECTIONS**  
39973 None.
- 39974 **SEE ALSO**  
39975 *getsockopt()*, *poll()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendmsg()*, *setsockopt()*,  
39976 *shutdown()*, *socket()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/socket.h> |
- 39977 **CHANGE HISTORY**  
39978 First released in Issue 6. Derived from the XNS, Issue 5.2 specification. |
- 39979 The wording of the mandatory [ELOOP] error condition is updated, and a second optional |  
39980 [ELOOP] error condition is added. |

39981 **NAME**

39982           setbuf — assign buffering to a stream

39983 **SYNOPSIS**

39984           #include &lt;stdio.h&gt;

39985           void setbuf(FILE \*restrict *stream*, char \*restrict *buf*);39986 **DESCRIPTION**

39987 *CX*       The functionality described on this reference page is aligned with the ISO C standard. Any  
39988       conflict between the requirements described here and the ISO C standard is unintentional. This  
39989       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

39990       Except that it returns no value, the function call:

39991           setbuf(*stream*, *buf*)

39992       shall be equivalent to:

39993           setvbuf(*stream*, *buf*, \_IOFBF, BUFSIZ)39994       if *buf* is not a null pointer, or to:39995           setvbuf(*stream*, *buf*, \_IONBF, BUFSIZ)39996       if *buf* is a null pointer.39997 **RETURN VALUE**39998       The *setbuf()* function shall return no value.39999 **ERRORS**

40000       No errors are defined.

40001 **EXAMPLES**

40002       None.

40003 **APPLICATION USAGE**

40004       A common source of error is allocating buffer space as an “automatic” variable in a code block,  
40005       and then failing to close the stream in the same block.

40006       With *setbuf()*, allocating a buffer of {BUFSIZ} bytes does not necessarily imply that all of  
40007       {BUFSIZ} bytes are used for the buffer area.

40008 **RATIONALE**

40009       None.

40010 **FUTURE DIRECTIONS**

40011       None.

40012 **SEE ALSO**40013       *fopen()*, *setvbuf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>40014 **CHANGE HISTORY**

40015       First released in Issue 1. Derived from Issue 1 of the SVID.

40016 **Issue 6**40017       The prototype for *setbuf()* is updated for alignment with the ISO/IEC 9899:1999 standard.

40018 **NAME**

40019 setcontext — set current user context

40020 **SYNOPSIS**

40021 xSI #include &lt;ucontext.h&gt;

40022 int setcontext(const ucontext\_t \*ucp);

40023

40024 **DESCRIPTION**40025 Refer to *getcontext()*.

40026 **NAME**

40027           setegid — set effective group ID

40028 **SYNOPSIS**

40029           #include &lt;unistd.h&gt;

40030           int setegid(gid\_t *gid*);40031 **DESCRIPTION**40032           If *gid* is equal to the real group ID or the saved set-group-ID, or if the process has appropriate  
40033           privileges, *setegid()* shall set the effective group ID of the calling process to *gid*; the real group  
40034           ID, saved set-group-ID, and any supplementary group IDs shall remain unchanged.40035           The *setegid()* function shall not affect the supplementary group list in any way.40036 **RETURN VALUE**40037           Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to  
40038           indicate the error.40039 **ERRORS**40040           The *setegid()* function shall fail if:40041           [EINVAL]           The value of the *gid* argument is invalid and is not supported by the  
40042           implementation.40043           [EPERM]           The process does not have appropriate privileges and *gid* does not match the  
40044           real group ID or the saved set-group-ID.40045 **EXAMPLES**

40046           None.

40047 **APPLICATION USAGE**

40048           None.

40049 **RATIONALE**40050           Refer to the RATIONALE section in *setuid()*.40051 **FUTURE DIRECTIONS**

40052           None.

40053 **SEE ALSO**40054           *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the  
40055           Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>40056 **CHANGE HISTORY**

40057           First released in Issue 6. Derived from the IEEE P1003.1a draft standard.

40058 **NAME**

40059        setenv — add or change environment variable

40060 **SYNOPSIS**

40061        #include &lt;stdlib.h&gt;

40062        int setenv(const char \*envname, const char \*envval, int overwrite);

40063 **DESCRIPTION**

40064        The *setenv()* function updates or adds a variable in the environment of the calling process. The *envname* argument points to a string containing the name of an environment variable to be added or altered. The environment variable shall be set to the value to which *envval* points. The function shall fail if *envname* points to a string which contains an '=' character. If the environment variable named by *envname* already exists and the value of *overwrite* is non-zero, the function shall return success and the environment shall be updated. If the environment variable named by *envname* already exists and the value of *overwrite* is zero, the function shall return success and the environment shall remain unchanged.

40072        If the application modifies *environ* or the pointers to which it points, the behavior of *setenv()* is undefined. The *setenv()* function shall update the list of pointers to which *environ* points.

40074        The strings described by *envname* and *envval* are copied by this function.

40075        The *setenv()* function need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

40077 **RETURN VALUE**

40078        Upon successful completion, zero shall be returned. Otherwise, -1 shall be returned, *errno* set to indicate the error, and the environment shall be unchanged.

40080 **ERRORS**

40081        The *setenv()* function shall fail if:

40082        [EINVAL]        The *name* argument is a null pointer, points to an empty string, or points to a string containing an '=' character.

40084        [ENOMEM]       Insufficient memory was available to add a variable or its value to the environment.

40086 **EXAMPLES**

40087        None.

40088 **APPLICATION USAGE**

40089        None.

40090 **RATIONALE**

40091        Unanticipated results may occur if *setenv()* changes the external variable *environ*. In particular, if the optional *envp* argument to *main()* is present, it is not changed, and thus may point to an obsolete copy of the environment (as may any other copy of *environ*). However, other than the aforementioned restriction, the developers of IEEE Std. 1003.1-200x intended that the traditional method of walking through the environment by way of the *environ* pointer must be supported.

40096        It was decided that *setenv()* should be required by this revision because it addresses a piece of missing functionality, and does not impose a significant burden on the implementor.

40098        There was considerable debate as to whether the System V *putenv()* function or the BSD *setenv()* function should be required as a mandatory function. The *setenv()* function was chosen because it permitted the implementation of *unsetenv()* function to delete environmental variables, without specifying an additional interface. The *putenv()* function is available as an XSI extension.

40103           The standard developers considered requiring that *setenv()* indicate an error when a call to it  
40104           would result in exceeding {ARG\_MAX}. The requirement was rejected since the condition might  
40105           be temporary, with the application eventually reducing the environment size. The ultimate  
40106           success or failure depends on the size at the time of a call to *exec*, which returns an indication of  
40107           this error condition.

40108 **FUTURE DIRECTIONS**

40109           None.

40110 **SEE ALSO**

40111           *getenv()*, *unsetenv()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>`, |  
40112           `<sys/types.h>`, `<unistd.h>`

40113 **CHANGE HISTORY**

40114           First released in Issue 6. Derived from the IEEE P1003.1a draft standard.

40115 **NAME**

40116            seteuid — set effective user ID

40117 **SYNOPSIS**

40118            #include &lt;unistd.h&gt;

40119            int seteuid(uid\_t uid);

40120 **DESCRIPTION**

40121            If *uid* is equal to the real user ID or the saved set-user-ID, or if the process has appropriate  
40122            privileges, *seteuid()* shall set the effective user ID of the calling process to *uid*; the real user ID  
40123            and saved set-user-ID shall remain unchanged.

40124            The *seteuid()* function shall not affect the supplementary group list in any way.

40125 **RETURN VALUE**

40126            Upon successful completion, 0 shall be returned; otherwise, -1 shall be returned and *errno* set to  
40127            indicate the error.

40128 **ERRORS**

40129            The *seteuid()* function shall fail if:

40130            [EINVAL]            The value of the *uid* argument is invalid and is not supported by the  
40131            implementation.

40132            [EPERM]            The process does not have appropriate privileges and *uid* does not match the  
40133            real group ID or the saved set-group-ID.

40134 **EXAMPLES**

40135            None.

40136 **APPLICATION USAGE**

40137            None.

40138 **RATIONALE**

40139            Refer to the RATIONALE section in *setuid()*.

40140 **FUTURE DIRECTIONS**

40141            None.

40142 **SEE ALSO**

40143            *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *setgid()*, *setregid()*, *setreuid()*, *setuid()*, the  
40144            Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>

40145 **CHANGE HISTORY**

40146            First released in Issue 6. Derived from the IEEE P1003.1a draft standard.

40147 **NAME**

40148        setgid — set-group-ID

40149 **SYNOPSIS**

40150        #include &lt;unistd.h&gt;

40151        int setgid(gid\_t *gid*);40152 **DESCRIPTION**40153        If the process has appropriate privileges, *setgid()* shall set the real group ID, effective group ID,  
40154        and the saved set-group-ID of the calling process to *gid*.40155        If the process does not have appropriate privileges, but *gid* is equal to the real group ID or the  
40156        saved set-group-ID, *setgid()* shall set the effective group ID to *gid*; the real group ID and saved  
40157        set-group-ID shall remain unchanged.40158        The *setgid()* function shall not affect the supplementary group list in any way.

40159        Any supplementary group IDs of the calling process shall remain unchanged.

40160 **RETURN VALUE**40161        Upon successful completion, 0 is returned. Otherwise, -1 shall be returned and *errno* set to  
40162        indicate the error.40163 **ERRORS**40164        The *setgid()* function shall fail if:40165        [EINVAL]        The value of the *gid* argument is invalid and is not supported by the  
40166        implementation.40167        [EPERM]        The process does not have appropriate privileges and *gid* does not match the  
40168        real group ID or the saved set-group-ID.40169 **EXAMPLES**

40170        None.

40171 **APPLICATION USAGE**

40172        None.

40173 **RATIONALE**40174        Refer to the RATIONALE section in *setuid()*.40175 **FUTURE DIRECTIONS**

40176        None.

40177 **SEE ALSO**40178        *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setregid()*, *setreuid()*, *setuid()*, the  
40179        Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>40180 **CHANGE HISTORY**

40181        First released in Issue 1. Derived from Issue 1 of the SVID.

40182 **Issue 4**40183        The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
40184        XSI-conformant systems.

40185        The following change is incorporated for alignment with the FIPS requirements:

- 40186
- All references to the saved set-user-ID are marked as extensions. This is because Issue 4  
40187        defines this mechanism as mandatory, whereas the ISO POSIX-1 standard defines that it is  
40188        only supported if `_POSIX_SAVED_IDS` is set.

40189 **Issue 6**

40190 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

40191 The following new requirements on POSIX implementations derive from alignment with the  
40192 Single UNIX Specification:

40193 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
40194 required for conforming implementations of previous POSIX specifications, it was not  
40195 required for UNIX applications.

40196 • Functionality associated with `_POSIX_SAVED_IDS` is now mandated. This is a FIPS  
40197 requirement.

40198 The following changes were made to align with the IEEE P1003.1a draft standard:

40199 • The effects of `setgid()` in processes without appropriate privileges are changed

40200 • A requirement that the supplementary group list is not affected is added.

40201 **NAME**

40202           setgrent — reset group database to first entry

40203 **SYNOPSIS**

40204 xSI       #include <grp.h>

40205           void setgrent(void);

40206

40207 **DESCRIPTION**

40208           Refer to *endgrent()*.

40209 **NAME**

40210           sethostent — network host database functions

40211 **SYNOPSIS**

40212           #include <netdb.h>

40213           void sethostent(int *stayopen*);

40214 **DESCRIPTION**

40215           Refer to *endhostent()*.

40216 **NAME**

40217       setitimer — set value of interval timer

40218 **SYNOPSIS**

40219 XSI       #include <sys/time.h>

40220       int setitimer(int *which*, const struct itimerval \*restrict *value*,  
40221                    struct itimerval \*restrict *ovalue*);

40222

40223 **DESCRIPTION**

40224       Refer to *getitimer()*.

40225 **NAME**

40226 setjmp — set jump point for a non-local goto

40227 **SYNOPSIS**

40228 #include &lt;setjmp.h&gt;

40229 int setjmp(jmp\_buf env);

40230 **DESCRIPTION**

40231 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
40232 conflict between the requirements described here and the ISO C standard is unintentional. This  
40233 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

40234 A call to *setjmp()*, shall save the calling environment in its *env* argument for later use by  
40235 *longjmp()*.

40236 It is unspecified whether *setjmp()* is a macro or a function. If a macro definition is suppressed in  
40237 order to access an actual function, or a program defines an external identifier with the name  
40238 *setjmp*, the behavior is undefined.

40239 All accessible objects have values as of the time *longjmp()* was called, except that the values of  
40240 objects of automatic storage duration which are local to the function containing the invocation of  
40241 the corresponding *setjmp()* which do not have volatile-qualified type and which are changed  
40242 between the *setjmp()* invocation and *longjmp()* call are indeterminate.

40243 An application shall ensure that an invocation of *setjmp()* appears in one of the following  
40244 contexts only:

- 40245 • The entire controlling expression of a selection or iteration statement
- 40246 • One operand of a relational or equality operator with the other operand an integral constant  
40247 expression, with the resulting expression being the entire controlling expression of a  
40248 selection or iteration statement
- 40249 • The operand of a unary '!' operator with the resulting expression being the entire  
40250 controlling expression of a selection or iteration
- 40251 • The entire expression of an expression statement (possibly cast to **void**)

40252 If the invocation appears in any other context, the behavior is undefined.

40253 **RETURN VALUE**

40254 If the return is from a direct invocation, *setjmp()* shall return 0. If the return is from a call to  
40255 *longjmp()*, *setjmp()* shall return a non-zero value.

40256 **ERRORS**

40257 No errors are defined.

40258 **EXAMPLES**

40259 None.

40260 **APPLICATION USAGE**

40261 In general, *sigsetjmp()* is more useful in dealing with errors and interrupts encountered in a low-  
40262 level subroutine of a program.

40263 **RATIONALE**

40264 None.

40265 **FUTURE DIRECTIONS**

40266 None.

40267 **SEE ALSO**40268 *longjmp()*, *sigsetjmp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <setjmp.h>40269 **CHANGE HISTORY**

40270 First released in Issue 1. Derived from Issue 1 of the SVID.

40271 **Issue 4**40272 This issue states that *setjmp()* is a macro or a function; previous issues stated that it was a macro.40273 Warnings have also been added about the suppression of a *setjmp()* macro definition.40274 Text describing the accessibility of objects after a *longjmp()* call is added to the DESCRIPTION.40275 This text is imported from the entry for *longjmp()*.40276 Text describing the contexts in which calls to *setjmp()* are valid is moved to the DESCRIPTION

40277 from the APPLICATION USAGE section.

40278 The APPLICATION USAGE section is changed to refer to *sigsetjmp()*.40279 **Issue 6**

40280 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

40281 **NAME**40282 setkey — set encoding key (**CRYPT**)40283 **SYNOPSIS**

40284 XSI #include &lt;stdlib.h&gt;

40285 void setkey(const char \*key);

40286

40287 **DESCRIPTION**

40288 The *setkey()* function provides (rather primitive) access to an implementation-defined encoding  
40289 algorithm. The argument of *setkey()* is an array of length 64 bytes containing only the bytes with  
40290 numerical value of 0 and 1. If this string is divided into groups of 8, the low-order bit in each  
40291 group is ignored; this gives a 56-bit key which is used by the algorithm. This is the key that shall  
40292 be used with the algorithm to encode a string *block* passed to *encrypt()*.

40293 The *setkey()* function shall not change the setting of *errno* if successful. An application wishing to  
40294 check for error situations should set *errno* to 0 before calling *setkey()*. If *errno* is non-zero on  
40295 return, an error has occurred.

40296 The *setkey()* function need not be reentrant. A function that is not required to be reentrant is not  
40297 required to be thread-safe.

40298 **RETURN VALUE**

40299 No values are returned.

40300 **ERRORS**40301 The *setkey()* function shall fail if:

40302 [ENOSYS] The functionality is not supported on this implementation.

40303 **EXAMPLES**

40304 None.

40305 **APPLICATION USAGE**

40306 Decoding need not be implemented in all environments. This is related to U.S. Government  
40307 restrictions on encryption and decryption routines: the DES decryption algorithm cannot be  
40308 exported outside the U.S. Historical practice has been to ship a different version of the  
40309 encryption library without the decryption feature in the routines supplied. Thus the exported  
40310 version of *encrypt()* does encoding but not decoding.

40311 **RATIONALE**

40312 None.

40313 **FUTURE DIRECTIONS**

40314 None.

40315 **SEE ALSO**40316 *crypt()*, *encrypt()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>40317 **CHANGE HISTORY**

40318 First released in Issue 1. Derived from Issue 1 of the SVID.

40319 **Issue 4**40320 The type of argument *key* is changed from **char\*** to **const char\***.

40321 The description of the array is put in terms of bytes instead of characters.

40322 The APPLICATION USAGE section is added.

40323 **Issue 5**

40324

The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

40325 **NAME**

40326 setlocale — set program locale

40327 **SYNOPSIS**

40328 #include &lt;locale.h&gt;

40329 char \*setlocale(int *category*, const char \**locale*);40330 **DESCRIPTION**

40331 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 40332 conflict between the requirements described here and the ISO C standard is unintentional. This  
 40333 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

40334 The *setlocale()* function selects the appropriate piece of the program's locale, as specified by the  
 40335 *category* and *locale* arguments, and may be used to change or query the program's entire locale or  
 40336 portions thereof. The value *LC\_ALL* for *category* names the program's entire locale; other values  
 40337 for *category* name only a part of the program's locale:

40338 *LC\_COLLATE* Affects the behavior of regular expressions and the collation functions.

40339 *LC\_CTYPE* Affects the behavior of regular expressions, character classification, character  
 40340 conversion functions, and wide-character functions.

40341 CX *LC\_MESSAGES* Affects what strings are expected by commands and utilities as affirmative or  
 40342 negative responses.

40343 XSI It also affects what strings are given by commands and utilities as affirmative  
 40344 or negative responses, and the content of messages.

40345 *LC\_MONETARY* Affects the behavior of functions that handle monetary values.

40346 *LC\_NUMERIC* Affects the behavior of functions that handle numeric values.

40347 *LC\_TIME* Affects the behavior of the time conversion functions.

40348 The *locale* argument is a pointer to a character string containing the required setting of *category*.  
 40349 The contents of this string are implementation-defined. In addition, the following preset values  
 40350 of *locale* are defined for all settings of *category*:

40351 CX "POSIX" Specifies the minimal environment for C-language translation called POSIX  
 40352 locale. If *setlocale()* is not invoked, the POSIX locale is the default at entry to  
 40353 *main()*.

40354 "C" Same as "POSIX".

40355 "" Specifies an implementation-defined native environment. For XSI-conformant  
 40356 systems, this corresponds to the value of the associated environment  
 40357 variables, *LC\_\** and *LANG*; see the Base Definitions volume of  
 40358 IEEE Std. 1003.1-200x, Chapter 7, Locale and the Base Definitions volume of  
 40359 IEEE Std. 1003.1-200x, Chapter 8, Environment Variables.

40360 A null pointer Used to direct *setlocale()* to query the current internationalized environment  
 40361 and return the name of the *locale()*.

40362 THR The locale state is common to all threads within a process.

40363 **RETURN VALUE**

40364 Upon successful completion, *setlocale()* shall return the string associated with the specified  
 40365 category for the new locale. Otherwise, *setlocale()* shall return a null pointer and the program's  
 40366 locale is not changed.

40367 A null pointer for *locale* causes *setlocale()* to return a pointer to the string associated with the  
40368 *category* for the program's current locale. The program's locale shall not be changed.

40369 The string returned by *setlocale()* is such that a subsequent call with that string and its associated  
40370 *category* shall restore that part of the program's locale. The application shall not modify the string  
40371 returned which may be overwritten by a subsequent call to *setlocale()*.

#### 40372 ERRORS

40373 No errors are defined.

#### 40374 EXAMPLES

40375 None.

#### 40376 APPLICATION USAGE

40377 The following code illustrates how a program can initialize the international environment for  
40378 one language, while selectively modifying the program's locale such that regular expressions  
40379 and string operations can be applied to text recorded in a different language:

```
40380 setlocale(LC_ALL, "De");
40381 setlocale(LC_COLLATE, "Fr@dict");
```

40382 Internationalized programs must call *setlocale()* to initiate a specific language operation. This can  
40383 be done by calling *setlocale()* as follows:

```
40384 setlocale(LC_ALL, "");
```

40385 Changing the setting of *LC\_MESSAGES* has no effect on catalogs that have already been opened  
40386 by calls to *catopen()*.

#### 40387 RATIONALE

40388 The ISO C standard defines a collection of functions to support internationalization. One of the  
40389 most significant aspects of these functions is a facility to set and query the *international*  
40390 *environment*. The international environment is a repository of information that affects the  
40391 behavior of certain functionality, namely:

- 40392 1. Character handling
- 40393 2. String handling (that is, collating)
- 40394 3. Date/time formatting
- 40395 4. Numeric editing

40396 The *setlocale()* function provides the application developer with the ability to set all or portions,  
40397 called *categories*, of the international environment. These categories correspond to the areas of  
40398 functionality, mentioned above. The syntax for *setlocale()* is as follows:

```
40399 char *setlocale(int category, const char *locale);
```

40400 where *category* is the name of one of five categories, namely:

```
40401 LC_COLLATE
40402 LC_CTYPE
40403 LC_MESSAGES
40404 LC_MONETARY
40405 LC_NUMERIC
40406 LC_TIME
```

40407 In addition, a special value called *LC\_ALL* directs *setlocale()* to set all categories.

40408 There are two primary uses of *setlocale()*:

- 40409           1. Querying the international environment to find out what it is set to  
 40410           2. Setting the international environment, or *locale*, to a specific value

40411           The behavior of *setlocale()* in these two areas is described below. Since it is difficult to describe  
 40412           the behavior in words, examples are used to illustrate the behavior of specific uses.

40413           To query the international environment, *setlocale()* is invoked with a specific category and the  
 40414           NULL pointer as the locale. The NULL pointer is a special directive to *setlocale()* that tells it to  
 40415           query rather than set the international environment. The following syntax is used to query the  
 40416           name of the international environment:

```
40417 setlocale({LC_ALL, LC_COLLATE, LC_CTYPE, LC_MESSAGES, LC_MONETARY, \
40418 LC_NUMERIC, LC_TIME}, (char *) NULL);
```

40419           The *setlocale()* function shall return the string corresponding to the current international  
 40420           environment. This value may be used by a subsequent call to *setlocale()* to reset the international  
 40421           environment to this value. However, it should be noted that the return value from *setlocale()* is a  
 40422           pointer to a static area within the function and is not guaranteed to remain unchanged (that is, it  
 40423           may be modified by a subsequent call to *setlocale()*). Therefore, if the purpose of calling  
 40424           *setlocale()* is to save the value of the current international environment so it can be changed and  
 40425           reset later, the return value should be copied to an array of **char** in the calling program.

40426           There are three ways to set the international environment with *setlocale()*:

40427           *setlocale(category, string)*

40428           This usage sets a specific *category* in the international environment to a specific value  
 40429           corresponding to the value of the *string*. A specific example is provided below:

```
40430 setlocale(LC_ALL, "Fr_FR.8859");
```

40431           In this example, all categories of the international environment are set to the locale  
 40432           corresponding to the string "Fr\_FR.8859", or to the French language as spoken in France  
 40433           using the ISO/IEC 8859-1:1998 standard codeset.

40434           If the string does not correspond to a valid locale, *setlocale()* shall return a NULL pointer  
 40435           and the international environment is not changed. Otherwise, *setlocale()* shall return the  
 40436           name of the locale just set.

40437           *setlocale(category, "C")*

40438           The ISO C standard states that one locale must exist on all conforming implementations.  
 40439           The name of the locale is C and corresponds to a minimal international environment needed  
 40440           to support the C programming language.

40441           *setlocale(category, "")*

40442           This sets a specific category to an implementation-defined default. For POSIX-conforming  
 40443           systems, this corresponds to the value of the environment variables.

#### 40444 FUTURE DIRECTIONS

40445           None.

#### 40446 SEE ALSO

40447           *exec*, *isalnum()*, *isalpha()*, *iscntrl()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*,  
 40448           *iswalnum()*, *iswalpha()*, *iswcntrl()*, *iswctype()*, *iswgraph()*, *iswlower()*, *iswprint()*, *iswpunct()*,  
 40449           *iswspace()*, *iswupper()*, *localeconv()*, *mblen()*, *mbstowcs()*, *mbtowc()*, *nl\_langinfo()*, *printf()*, *scanf()*,  
 40450           *setlocale()*, *strcoll()*, *strerror()*, *strfmon()*, *strtod()*, *strxfrm()*, *tolower()*, *toupper()*, *towlower()*,  
 40451           *towupper()*, *wscoll()*, *wcstod()*, *wcstombs()*, *wcsxfrm()*, *wctomb()*, the Base Definitions volume of  
 40452           IEEE Std. 1003.1-200x, <**langinfo.h**>, <**locale.h**>

40453 **CHANGE HISTORY**

40454 First released in Issue 3.

40455 **Issue 4**

40456 The description of *LC\_MESSAGES* is extended to indicate that this category also determines  
40457 what strings are produced by commands and utilities for affirmative and negative responses,  
40458 and that it affects the content of other program messages. This is marked as an extension.

40459 References to *nl\_langinfo()* are removed.

40460 The description of the implementation-defined native locale (" ") is clarified by stating the related  
40461 environment variables explicitly.

40462 The APPLICATION USAGE section is expanded.

40463 The following changes are incorporated for alignment with the ISO C standard and the  
40464 ISO POSIX-1 standard:

- 40465 • The type of the argument *locale* is changed from **char\*** to **const char\***.

- 40466 • The name "POSIX" is added to the list of standard locale names.

40467 The following change is incorporated for alignment with the ISO POSIX-2 standard:

- 40468 • The *LC\_MESSAGES* value for *category* is added to the DESCRIPTION.

40469 **Issue 5**

40470 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

40471 **Issue 6**

40472 Extensions beyond the ISO C standard are now marked.

40473 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

40474 **NAME**

40475 setlogmask — set log priority mask

40476 **SYNOPSIS**

40477 xSI #include &lt;syslog.h&gt;

40478 int setlogmask(int *maskpri*);

40479

40480 **DESCRIPTION**40481 Refer to *closelog()*.

# setnetent()

40482 **NAME**

40483       setnetent — network database function

40484 **SYNOPSIS**

40485       #include <netdb.h>

40486       void setnetent(int stayopen);

40487 **DESCRIPTION**

40488       Refer to *endnetent()*.

40489 **NAME**

40490 setpgid — set process group ID for job control

40491 **SYNOPSIS**

40492 #include &lt;unistd.h&gt;

40493 int setpgid(pid\_t pid, pid\_t pgid);

40494 **DESCRIPTION**

40495 The *setpgid()* function is used either to join an existing process group or create a new process  
 40496 group within the session of the calling process. The process group ID of a session leader shall  
 40497 not change. Upon successful completion, the process group ID of the process with a process ID  
 40498 that matches *pid* shall be set to *pgid*. As a special case, if *pid* is 0, the process ID of the calling  
 40499 process shall be used. Also, if *pgid* is 0, the process group ID of the indicated process shall be  
 40500 used.

40501 **RETURN VALUE**

40502 Upon successful completion, *setpgid()* shall return 0; otherwise, -1 shall be returned and *errno*  
 40503 shall be set to indicate the error.

40504 **ERRORS**40505 The *setpgid()* function shall fail if:

40506 [EACCES] The value of the *pid* argument matches the process ID of a child process of the  
 40507 calling process and the child process has successfully executed one of the *exec*  
 40508 functions.

40509 [EINVAL] The value of the *pgid* argument is less than 0, or is not a value supported by  
 40510 the implementation.

40511 [EPERM] The process indicated by the *pid* argument is a session leader.

40512 [EPERM] The value of the *pid* argument matches the process ID of a child process of the  
 40513 calling process and the child process is not in the same session as the calling  
 40514 process.

40515 [EPERM] The value of the *pgid* argument is valid but does not match the process ID of  
 40516 the process indicated by the *pid* argument and there is no process with a  
 40517 process group ID that matches the value of the *pgid* argument in the same  
 40518 session as the calling process.

40519 [ESRCH] The value of the *pid* argument does not match the process ID of the calling  
 40520 process or of a child process of the calling process.

40521 **EXAMPLES**

40522 None.

40523 **APPLICATION USAGE**

40524 None.

40525 **RATIONALE**

40526 The *setpgid()* function is used to group processes together for the purpose of signaling,  
 40527 placement in foreground or background, and other job control actions.

40528 The *setpgid()* function is similar to the *setpgrp()* function of 4.2 BSD, except that 4.2 BSD allowed  
 40529 the specified new process group to assume any value. This presents certain security problems  
 40530 and is more flexible than necessary to support job control.

40531 To provide tighter security, *setpgid()* only allows the calling process to join a process group  
 40532 already in use inside its session or create a new process group whose process group ID was

40533 equal to its process ID.

40534 When a job control shell spawns a new job, the processes in the job must be placed into a new  
40535 process group via *setpgid()*. There are two timing constraints involved in this action:

40536 1. The new process must be placed in the new process group before the appropriate program  
40537 is launched via one of the *exec* functions.

40538 2. The new process must be placed in the new process group before the shell can correctly  
40539 send signals to the new process group.

40540 To address these constraints, the following actions are performed. The new processes call  
40541 *setpgid()* to alter their own process groups after *fork()* but before *exec*. This satisfies the first  
40542 constraint. Under 4.3 BSD, the second constraint is satisfied by the synchronization property of  
40543 *vfork()*; that is, the shell is suspended until the child has completed the *exec*, thus ensuring that  
40544 the child has completed the *setpgid()*. A new version of *fork()* with this same synchronization  
40545 property was considered, but it was decided instead to merely allow the parent shell process to  
40546 adjust the process group of its child processes via *setpgid()*. Both timing constraints are now  
40547 satisfied by having both the parent shell and the child attempt to adjust the process group of the  
40548 child process; it does not matter which succeeds first.

40549 Because it would be confusing to an application to have its process group change after it began  
40550 executing (that is, after *exec*), and because the child process would already have adjusted its  
40551 process group before this, the [EACCES] error was added to disallow this.

40552 One non-obvious use of *setpgid()* is to allow a job control shell to return itself to its original  
40553 process group (the one in effect when the job control shell was executed). A job control shell  
40554 does this before returning control back to its parent when it is terminating or suspending itself as  
40555 a way of restoring its job control “state” back to what its parent would expect. (Note that the  
40556 original process group of the job control shell typically matches the process group of its parent,  
40557 but this is not necessarily always the case.)

#### 40558 FUTURE DIRECTIONS

40559 None.

#### 40560 SEE ALSO

40561 *exec*, *getpgrp()*, *setsid()*, *tcsetpgrp()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
40562 <sys/types.h>, <unistd.h>

#### 40563 CHANGE HISTORY

40564 First released in Issue 3.

40565 Entry included for alignment with the POSIX.1-1988 standard.

#### 40566 Issue 4

40567 The function is no longer marked as OPTIONAL FUNCTIONALITY.

40568 The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
40569 XSI-conformant systems.

40570 The <unistd.h> header is added to the SYNOPSIS section.

40571 The DESCRIPTION in Issue 3 defined the behavior of this function for implementations that  
40572 either supported or did not support job control. As job control is defined as mandatory in Issue  
40573 4, only the former of these is now described.

40574 The [ENOSYS] error is removed from the ERRORS section.

40575 **Issue 6**

40576 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

40577 The following new requirements on POSIX implementations derive from alignment with the  
40578 Single UNIX Specification:

- 40579 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
40580 required for conforming implementations of previous POSIX specifications, it was not  
40581 required for UNIX applications.
- 40582 • The `setpgid()` function is mandatory since `_POSIX_JOB_CONTROL` is required to be defined  
40583 in this issue. This is a FIPS requirement.

40584 **NAME**

40585           setpgrp — set process group ID

40586 **SYNOPSIS**

40587 XSI       #include &lt;unistd.h&gt;

40588           pid\_t setpgrp(void);

40589

40590 **DESCRIPTION**

40591           If the calling process is not already a session leader, *setpgrp()* sets the process group ID of the  
40592           calling process to the process ID of the calling process. If *setpgrp()* creates a new session, then  
40593           the new session has no controlling terminal.

40594           The *setpgrp()* function has no effect when the calling process is a session leader.

40595 **RETURN VALUE**40596           Upon completion, *setpgrp()* shall return the process group ID.40597 **ERRORS**

40598           No errors are defined.

40599 **EXAMPLES**

40600           None.

40601 **APPLICATION USAGE**

40602           None.

40603 **RATIONALE**

40604           None.

40605 **FUTURE DIRECTIONS**

40606           None.

40607 **SEE ALSO**

40608           *exec*, *fork()*, *getpid()*, *getsid()*, *kill()*, *setpgid()*, *setsid()*, the Base Definitions volume of  
40609           IEEE Std. 1003.1-200x, <unistd.h>

40610 **CHANGE HISTORY**

40611           First released in Issue 4, Version 2.

40612 **Issue 5**

40613           Moved from X/OPEN UNIX extension to BASE.

40614 **NAME**

40615           setpriority — set the nice value

40616 **SYNOPSIS**

40617 xSI       #include &lt;sys/resource.h&gt;

40618           int setpriority(int *which*, id\_t *who*, int *nice*);

40619

40620 **DESCRIPTION**40621           Refer to *getpriority()*.

# setprotoent()

40622 **NAME**

40623           setprotoent — network protocol database functions

40624 **SYNOPSIS**

40625           #include <netdb.h>

40626           void setprotoent(int *stayopen*);

40627 **DESCRIPTION**

40628           Refer to *endprotoent()*.

40629 **NAME**

40630           setpwent — user database function

40631 **SYNOPSIS**

40632 xSI       #include &lt;pwd.h&gt;

40633           void setpwent(void);

40634

40635 **DESCRIPTION**40636           Refer to *endpwent()*.

40637 **NAME**

40638 setregid — set real and effective group IDs

40639 **SYNOPSIS**

40640 XSI #include &lt;unistd.h&gt;

40641 int setregid(gid\_t rgid, gid\_t egid);

40642

40643 **DESCRIPTION**40644 The *setregid()* function is used to set the real and effective group IDs of the calling process.40645 If *rgid* is  $-1$ , the real group ID shall not be changed; if *egid* is  $-1$ , the effective group ID shall not  
40646 be changed.

40647 The real and effective group IDs may be set to different values in the same call.

40648 Only a process with appropriate privileges can set the real group ID and the effective group ID  
40649 to any valid value.40650 A non-privileged process can set either the real group ID to the saved set-group-ID from one of  
40651 the *exec* family of functions, or the effective group ID to the saved set-group-ID or the real group  
40652 ID.

40653 Any supplementary group IDs of the calling process remain unchanged.

40654 **RETURN VALUE**40655 Upon successful completion, 0 shall be returned. Otherwise,  $-1$  shall be returned and *errno* set to  
40656 indicate the error, and neither of the group IDs are changed.40657 **ERRORS**40658 The *setregid()* function shall fail if:40659 [EINVAL] The value of the *rgid* or *egid* argument is invalid or out-of-range.40660 [EPERM] The process does not have appropriate privileges and a change other than  
40661 changing the real group ID to the saved set-group-ID, or changing the  
40662 effective group ID to the real group ID or the saved set-group-ID, was  
40663 requested.40664 **EXAMPLES**

40665 None.

40666 **APPLICATION USAGE**40667 If a set-group-ID process sets its effective group ID to its real group ID, it can still set its effective  
40668 group ID back to the saved set-group-ID.40669 **RATIONALE**

40670 None.

40671 **FUTURE DIRECTIONS**

40672 None.

40673 **SEE ALSO**40674 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setreuid()*, *setuid()*, the  
40675 Base Definitions volume of IEEE Std. 1003.1-200x, <unistd.h>40676 **CHANGE HISTORY**

40677 First released in Issue 4, Version 2.

40678 **Issue 5**

40679 Moved from X/OPEN UNIX extension to BASE.

40680 The DESCRIPTION is updated to indicate that the saved set-group-ID can be set by any of the  
40681 *exec* family of functions, not just *execev()*.

40682 **NAME**

40683 setreuid — set real and effective user IDs

40684 **SYNOPSIS**

40685 XSI #include &lt;unistd.h&gt;

40686 int setreuid(uid\_t ruid, uid\_t euid);

40687

40688 **DESCRIPTION**

40689 The *setreuid()* function sets the real and effective user IDs of the current process to the values  
 40690 specified by the *ruid* and *euid* arguments. If *ruid* or *euid* is  $-1$ , the corresponding effective or real  
 40691 user ID of the current process is left unchanged.

40692 A process with appropriate privileges can set either ID to any value. An unprivileged process  
 40693 can only set the effective user ID if the *euid* argument is equal to either the real, effective, or  
 40694 saved user ID of the process.

40695 It is unspecified whether a process without appropriate privileges is permitted to change the real  
 40696 user ID to match the current real, effective, or saved set-user-ID of the process.

40697 **RETURN VALUE**

40698 Upon successful completion, 0 shall be returned. Otherwise,  $-1$  shall be returned and *errno* set to  
 40699 indicate the error.

40700 **ERRORS**40701 The *setreuid()* function shall fail if:

|       |          |                                                                                  |  |
|-------|----------|----------------------------------------------------------------------------------|--|
| 40702 | [EINVAL] | The value of the <i>ruid</i> or <i>euid</i> argument is invalid or out-of-range. |  |
| 40703 | [EPERM]  | The current process does not have appropriate privileges, and either an          |  |
| 40704 |          | attempt was made to change the effective user ID to a value other than the       |  |
| 40705 |          | real user ID or the saved set-user-ID or an attempt was made to change the       |  |
| 40706 |          | real user ID to a value not permitted by the implementation.                     |  |

40707 **EXAMPLES**40708 **Setting the Effective User ID to the Real User ID**

40709 The following example sets the effective user ID of the calling process to the real user ID, so that  
 40710 files created later will be owned by the current user.

```
40711 #include <unistd.h>
40712 #include <sys/types.h>
40713 ...
40714 setreuid(getuid(), getuid());
40715 ...
```

40716 **APPLICATION USAGE**

40717 None.

40718 **RATIONALE**

40719 None.

40720 **FUTURE DIRECTIONS**

40721 None.

40722 **SEE ALSO**

40723 *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setuid()*, the Base  
40724 Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

40725 **CHANGE HISTORY**

40726 First released in Issue 4, Version 2.

40727 **Issue 5**

40728 Moved from X/OPEN UNIX extension to BASE.

40729 **NAME**

40730 setrlimit — control maximum resource consumption

40731 **SYNOPSIS**

40732 XSI #include <sys/resource.h>

40733 int setrlimit(int resource, const struct rlimit \*rlp);

40734

40735 **DESCRIPTION**

40736 Refer to *getrlimit()*.

40737 **NAME**

40738           setservent — network services database functions

40739 **SYNOPSIS**

40740           #include <netdb.h>

40741           void setservent(int *stayopen*);

40742 **DESCRIPTION**

40743           Refer to *endservent()*.

40744 **NAME**

40745 setsid — create session and set process group ID

40746 **SYNOPSIS**

40747 #include &lt;unistd.h&gt;

40748 pid\_t setsid(void);

40749 **DESCRIPTION**

40750 The *setsid()* function creates a new session, if the calling process is not a process group leader.  
40751 Upon return the calling process shall be the session leader of this new session, shall be the  
40752 process group leader of a new process group, and shall have no controlling terminal. The  
40753 process group ID of the calling process shall be set equal to the process ID of the calling process.  
40754 The calling process shall be the only process in the new process group and the only process in  
40755 the new session.

40756 **RETURN VALUE**

40757 Upon successful completion, *setsid()* shall return the value of the new process group ID of the  
40758 calling process. Otherwise, it shall return (**pid\_t**)-1 and set *errno* to indicate the error.

40759 **ERRORS**40760 The *setsid()* function shall fail if:

40761 [EPERM] The calling process is already a process group leader, or the process group ID  
40762 of a process other than the calling process matches the process ID of the  
40763 calling process.

40764 **EXAMPLES**

40765 None.

40766 **APPLICATION USAGE**

40767 None.

40768 **RATIONALE**

40769 The *setsid()* function is similar to the *setpgrp()* function of System V. System V, without job  
40770 control, groups processes into process groups and creates new process groups via *setpgrp()*; only  
40771 one process group may be part of a login session.

40772 Job control allows multiple process groups within a login session. In order to limit job control  
40773 actions so that they can only affect processes in the same login session, this volume of  
40774 IEEE Std. 1003.1-200x adds the concept of a session that is created via *setsid()*. The *setsid()*  
40775 function also creates the initial process group contained in the session. Additional process  
40776 groups can be created via the *setpgid()* function. A System V process group would correspond to  
40777 a POSIX System Interfaces session containing a single POSIX process group. Note that this  
40778 function requires that the calling process not be a process group leader. The usual way to ensure  
40779 this is true is to create a new process with *fork()* and have it call *setsid()*. The *fork()* function  
40780 guarantees that the process ID of the new process does not match any existing process group ID.

40781 **FUTURE DIRECTIONS**

40782 None.

40783 **SEE ALSO**

40784 *getsid()*, *setpgid()*, *setpgrp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>,  
40785 <unistd.h>

40786 **CHANGE HISTORY**

40787 First released in Issue 3.

40788 Entry included for alignment with the POSIX.1-1988 standard.

40789 **Issue 4**40790 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on XSI-conformant systems.40792 The `<unistd.h>` header is added to the SYNOPSIS section.40793 The argument list is explicitly defined as **void**.40794 **Issue 6**40795 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

40796 The following new requirements on POSIX implementations derive from alignment with the Single UNIX Specification:

- 40798
- The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was required for conforming implementations of previous POSIX specifications, it was not required for UNIX applications.
- 40799
- 
- 40800

## 40801 NAME

40802 setsockopt — set the socket options

## 40803 SYNOPSIS

40804 #include &lt;sys/socket.h&gt;

```
40805 int setsockopt(int socket, int level, int option_name,
40806 const void *option_value, socklen_t option_len);
```

## 40807 DESCRIPTION

40808 The *setsockopt()* function sets the option specified by the *option\_name* argument, at the protocol  
 40809 level specified by the *level* argument, to the value pointed to by the *option\_value* argument for the  
 40810 socket associated with the file descriptor specified by the *socket* argument.

40811 The *level* argument specifies the protocol level at which the option resides. To set options at the  
 40812 socket level, specify the *level* argument as SOL\_SOCKET. To set options at other levels, supply  
 40813 the appropriate *level* identifier for the protocol controlling the option. For example, to indicate  
 40814 that an option is interpreted by the TCP (Transport Control Protocol), set *level* to IPPROTO\_TCP  
 40815 as defined in the <netinet/in.h> header.

40816 The *option\_name* argument specifies a single option to set. The *option\_name* argument and any  
 40817 specified options are passed uninterpreted to the appropriate protocol module for  
 40818 interpretations. The <sys/socket.h> header defines the socket-level options. The options are as  
 40819 follows:

40820 SO\_DEBUG Turns on recording of debugging information. This option enables or  
 40821 disables debugging in the underlying protocol modules. This option takes  
 40822 an **int** value. This is a Boolean option.

40823 SO\_BROADCAST Permits sending of broadcast messages, if this is supported by the  
 40824 protocol. This option takes an **int** value. This is a Boolean option.

40825 SO\_REUSEADDR Specifies that the rules used in validating addresses supplied to *bind()*  
 40826 should allow reuse of local addresses, if this is supported by the protocol.  
 40827 This option takes an **int** value. This is a Boolean option.

40828 SO\_KEEPALIVE Keeps connections active by enabling the periodic transmission of  
 40829 messages, if this is supported by the protocol. This option takes an **int**  
 40830 value.

40831 If the connected socket fails to respond to these messages, the connection  
 40832 is broken and threads writing to that socket are notified with a SIGPIPE  
 40833 signal.

40834 This is a Boolean option.

40835 SO\_LINGER Lingers on a *close()* if data is present. This option controls the action  
 40836 taken when unsent messages queue on a socket and *close()* is performed.  
 40837 If SO\_LINGER is set, the system blocks the process during *close()* until it  
 40838 can transmit the data or until the time expires. If SO\_LINGER is not  
 40839 specified, and *close()* is issued, the system handles the call in a way that  
 40840 allows the process to continue as quickly as possible. This option takes a  
 40841 **linger** structure, as defined in the <sys/socket.h> header, to specify the  
 40842 state of the option and linger interval.

40843 SO\_OOBINLINE Leaves received out-of-band data (data marked urgent) inline. This  
 40844 option takes an **int** value. This is a Boolean option.

|       |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 40845 | SO_SNDBUF           | Sets send buffer size. This option takes an <b>int</b> value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 40846 | SO_RCVBUF           | Sets receive buffer size. This option takes an <b>int</b> value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| 40847 | SO_DONTROUTE        | Requests that outgoing messages bypass the standard routing facilities. The destination shall be on a directly-connected network, and messages are directed to the appropriate network interface according to the destination address. The effect, if any, of this option depends on what protocol is in use. This option takes an <b>int</b> value. This is a Boolean option.                                                                                                                                                                                                                                                                                                                    |
| 40848 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40849 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40850 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40851 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40852 | SO_RCVLOWAT         | Sets the minimum number of bytes to process for socket input operations. The default value for SO_RCVLOWAT is 1. If SO_RCVLOWAT is set to a larger value, blocking receive calls normally wait until they have received the smaller of the low water mark value or the requested amount. (They may return less than the low water mark if an error occurs, a signal is caught, or the type of data next in the receive queue is different than that returned; for example, out-of-band data.) This option takes an <b>int</b> value. Note that not all implementations allow this option to be set.                                                                                               |
| 40853 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40854 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40855 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40856 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40857 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40858 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40859 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40860 | SO_RCVTIMEO         | Sets the timeout value that specifies the maximum amount of time an input function waits until it completes. It accepts a <b>timeval</b> structure with the number of seconds and microseconds specifying the limit on how long to wait for an input operation to complete. If a receive operation has blocked for this much time without receiving additional data, it shall return with a partial count or <i>errno</i> set to [EAGAIN] or [EWOULDBLOCK] if no data is received. The default for this option is zero, which indicates that a receive operation shall not time out. This option takes a <b>timeval</b> structure. Note that not all implementations allow this option to be set. |
| 40861 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40862 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40863 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40864 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40865 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40866 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40867 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40868 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40869 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40870 | SO_SNDLOWAT         | Sets the minimum number of bytes to process for socket output operations. Non-blocking output operations shall process no data if flow control does not allow the smaller of the send low water mark value or the entire request to be processed. This option takes an <b>int</b> value. Note that not all implementations allow this option to be set.                                                                                                                                                                                                                                                                                                                                           |
| 40871 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40872 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40873 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40874 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40875 | SO_SNDTIMEO         | Sets the timeout value specifying the amount of time that an output function blocks because flow control prevents data from being sent. If a send operation has blocked for this time, it shall return with a partial count or with <i>errno</i> set to [EAGAIN] or [EWOULDBLOCK] if no data is sent. The default for this option is zero, which indicates that a send operation shall not time out. This option stores a <b>timeval</b> structure. Note that not all implementations allow this option to be set.                                                                                                                                                                                |
| 40876 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40877 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40878 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40879 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40880 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40881 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40882 |                     | For Boolean options, 0 indicates that the option is disabled and 1 indicates that the option is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 40883 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40884 |                     | Options at other protocol levels vary in format and name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| 40885 | <b>RETURN VALUE</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40886 |                     | Upon successful completion, <i>setsockopt()</i> shall return 0. Otherwise, -1 shall be returned and <i>errno</i> set to indicate the error.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| 40887 |                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40888 | <b>ERRORS</b>       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40889 |                     | The <i>setsockopt()</i> function shall fail if:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 40890 | [EBADF]             | The <i>socket</i> argument is not a valid file descriptor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

|       |                          |                                                                                                                                                                                                                                                                                                                                                               |
|-------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 40891 | [EDOM]                   | The send and receive timeout values are too big to fit into the timeout fields in the socket structure.                                                                                                                                                                                                                                                       |
| 40892 |                          |                                                                                                                                                                                                                                                                                                                                                               |
| 40893 | [EINVAL]                 | The specified option is invalid at the specified socket level or the socket has been shut down.                                                                                                                                                                                                                                                               |
| 40894 |                          |                                                                                                                                                                                                                                                                                                                                                               |
| 40895 | [EISCONN]                | The socket is already connected, and a specified option cannot be set while the socket is connected.                                                                                                                                                                                                                                                          |
| 40896 |                          |                                                                                                                                                                                                                                                                                                                                                               |
| 40897 | [ENOPROTOPT]             |                                                                                                                                                                                                                                                                                                                                                               |
| 40898 |                          | The option is not supported by the protocol.                                                                                                                                                                                                                                                                                                                  |
| 40899 | [ENOTSOCK]               | The <i>socket</i> argument does not refer to a socket.                                                                                                                                                                                                                                                                                                        |
| 40900 |                          | The <i>setsockopt()</i> function may fail if:                                                                                                                                                                                                                                                                                                                 |
| 40901 | [ENOMEM]                 | There was insufficient memory available for the operation to complete.                                                                                                                                                                                                                                                                                        |
| 40902 | [ENOBUFS]                | Insufficient resources are available in the system to complete the call.                                                                                                                                                                                                                                                                                      |
| 40903 | <b>EXAMPLES</b>          |                                                                                                                                                                                                                                                                                                                                                               |
| 40904 |                          | None.                                                                                                                                                                                                                                                                                                                                                         |
| 40905 | <b>APPLICATION USAGE</b> |                                                                                                                                                                                                                                                                                                                                                               |
| 40906 |                          | The <i>setsockopt()</i> function provides an application program with the means to control socket behavior. An application program can use <i>setsockopt()</i> to allocate buffer space, control timeouts, or permit socket data broadcasts. The <code>&lt;sys/socket.h&gt;</code> header defines the socket-level options available to <i>setsockopt()</i> . |
| 40907 |                          |                                                                                                                                                                                                                                                                                                                                                               |
| 40908 |                          |                                                                                                                                                                                                                                                                                                                                                               |
| 40909 |                          |                                                                                                                                                                                                                                                                                                                                                               |
| 40910 |                          | Options may exist at multiple protocol levels. The SO_ options are always present at the uppermost socket level.                                                                                                                                                                                                                                              |
| 40911 |                          |                                                                                                                                                                                                                                                                                                                                                               |
| 40912 | <b>RATIONALE</b>         |                                                                                                                                                                                                                                                                                                                                                               |
| 40913 |                          | None.                                                                                                                                                                                                                                                                                                                                                         |
| 40914 | <b>FUTURE DIRECTIONS</b> |                                                                                                                                                                                                                                                                                                                                                               |
| 40915 |                          | None.                                                                                                                                                                                                                                                                                                                                                         |
| 40916 | <b>SEE ALSO</b>          |                                                                                                                                                                                                                                                                                                                                                               |
| 40917 |                          | Section 2.10 (on page 562), <i>bind()</i> , <i>endprotoent()</i> , <i>getsockopt()</i> , <i>socket()</i> , the Base Definitions                                                                                                                                                                                                                               |
| 40918 |                          | volume of IEEE Std. 1003.1-200x, <code>&lt;netinet/in.h&gt;</code> , <code>&lt;sys/socket.h&gt;</code>                                                                                                                                                                                                                                                        |
| 40919 | <b>CHANGE HISTORY</b>    |                                                                                                                                                                                                                                                                                                                                                               |
| 40920 |                          | First released in Issue 6. Derived from the XNS, Issue 5.2 specification.                                                                                                                                                                                                                                                                                     |

40921 **NAME**

40922            setstate — switch pseudorandom number generator state arrays

40923 **SYNOPSIS**

40924 xSI        #include &lt;stdlib.h&gt;

40925            char \*setstate(const char \*state);

40926

40927 **DESCRIPTION**40928            Refer to *initstate()*.

40929 **NAME**

40930           setuid — set user ID

40931 **SYNOPSIS**

40932           #include &lt;unistd.h&gt;

40933           int setuid(uid\_t uid);

40934 **DESCRIPTION**40935           If the process has appropriate privileges, *setuid()* shall set the real user ID, effective user ID, and  
40936           the saved set-user-ID of the calling process to *uid*.40937           If the process does not have appropriate privileges, but *uid* is equal to the real user ID or the  
40938           saved set-user-ID, *setuid()* shall set the effective user ID to *uid*; the real user ID and saved set-  
40939           user-ID shall remain unchanged.40940           The *setuid()* function shall not affect the supplementary group list in any way.40941 **RETURN VALUE**40942           Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
40943           indicate the error.40944 **ERRORS**40945           The *setuid()* function shall fail, return -1, and set *errno* to the corresponding value if one or more  
40946           of the following are true:40947           [EINVAL]           The value of the *uid* argument is invalid and not supported by the  
40948           implementation.40949           [EPERM]           The process does not have appropriate privileges and *uid* does not match the  
40950           real user ID or the saved set-user-ID.40951 **EXAMPLES**

40952           None.

40953 **APPLICATION USAGE**

40954           None.

40955 **RATIONALE**40956           The various behaviors of the *setuid()* and *setgid()* functions when called by non-privileged  
40957           processes reflect the behavior of different historical implementations. For portability, it is  
40958           recommended that new non-privileged applications use the *seteuid()* and *setegid()* functions  
40959           instead.40960           The saved set-user-ID capability allows a program to regain the effective user ID established at  
40961           the last *exec* call. Similarly, the saved set-group-ID capability allows a program to regain the  
40962           effective group ID established at the last *exec* call. These capabilities are derived from System V.  
40963           Without them, a program might have to run as superuser in order to perform the same  
40964           functions, because superuser can write on the user's files. This is a problem because such a  
40965           program can write on any user's files, and so must be carefully written to emulate the  
40966           permissions of the calling process properly. In System V, these capabilities have traditionally  
40967           been implemented only via the *setuid()* and *setgid()* functions for non-privileged processes. The  
40968           fact that the behavior of those functions was different for privileged processes made them  
40969           difficult to use. The POSIX.1-1990 standard defined the *setuid()* function to behave differently  
40970           for privileged and unprivileged users. When the caller had the appropriate privilege, the  
40971           function set the calling process's real user ID, effective user ID, and saved set-user ID on  
40972           implementations that supported it. When the caller did not have the appropriate privilege, the  
40973           function set only the effective user ID, subject to permission checks. The former use is generally  
40974           needed for utilities like *login* and *su*, which are not portable applications and thus outside the

40975 scope of IEEE Std. 1003.1-200x. These utilities wish to change the user ID irrevocably to a new  
 40976 value, generally that of an unprivileged user. The latter use is needed for portable applications  
 40977 that are installed with the set-user-ID bit and need to perform operations using the real user ID.

40978 IEEE Std. 1003.1-200x augments the latter functionality with a mandatory feature named  
 40979 `_POSIX_SAVED_IDS`. This feature permits a set-user-ID application to switch its effective user  
 40980 ID back and forth between the values of its *exec*-time real user ID and effective user ID.  
 40981 Unfortunately, the POSIX.1-1990 standard did not permit a portable application using this  
 40982 feature to work properly when it happened to be executed with the (implementation-defined)  
 40983 appropriate privilege. Furthermore, the application did not even have a means to tell whether it  
 40984 had this privilege. Because the saved set-user-ID feature is quite desirable for applications, as  
 40985 evidenced by the fact that NIST required it in FIPS 151-2, it has been mandated by  
 40986 IEEE Std. 1003.1-200x. However, there are implementors who have been reluctant to support it  
 40987 given the limitation described above.

40988 The 4.3BSD system handles the problem by supporting separate functions: *setuid()* (which  
 40989 always sets both the real and effective user IDs, like *setuid()* in IEEE Std. 1003.1-200x for  
 40990 privileged users), and *seteuid()* (which always sets just the effective user ID, like *setuid()* in  
 40991 IEEE Std. 1003.1-200x for non-privileged users). This separation of functionality into distinct  
 40992 functions seems desirable. 4.3BSD does not support the saved set-user-ID feature. It supports  
 40993 similar functionality of switching the effective user ID back and forth via *setreuid()*, which  
 40994 permits reversing the real and effective user IDs. This model seems less desirable than the saved  
 40995 set-user-ID because the real user ID changes as a side effect. The current 4.4BSD includes saved  
 40996 effective IDs and uses them for *seteuid()* and *setegid()* as described above. The *setreuid()* and  
 40997 *setregid()* functions will be deprecated or removed.

40998 The solution here is:

- 40999 • Require that all implementations support the functionality of the saved set-user-ID, which is  
 41000 set by the *exec* functions and by privileged calls to *setuid()*.
- 41001 • Add the *seteuid()* and *setegid()* functions as portable alternatives to *setuid()* and *setgid()* for  
 41002 non-privileged and privileged processes.

41003 Historical systems have provided two mechanisms for a set-user-ID process to change its  
 41004 effective user ID to be the same as its real user ID in such a way that it could return to the  
 41005 original effective user ID: the use of the *setuid()* function in the presence of a saved set-user-ID,  
 41006 or the use of the BSD *setreuid()* function, which was able to swap the real and effective user IDs.  
 41007 The changes included in IEEE Std. 1003.1-200x provide a new mechanism using *seteuid()*  
 41008 in conjunction with a saved set-user-ID. Thus, all implementations with the new *seteuid()*  
 41009 mechanism will have a saved set-user-ID for each process, and most of the behavior controlled  
 41010 by `_POSIX_SAVED_IDS` has been changed to agree with the case where the option was defined.  
 41011 The *kill()* function is an exception. Implementors of the new *seteuid()* mechanism will generally  
 41012 be required to maintain compatibility with the older mechanisms previously supported by their  
 41013 systems. However, compatibility with this use of *setreuid()* and with the `_POSIX_SAVED_IDS`  
 41014 behavior of *kill()* is unfortunately complicated. If an implementation with a saved set-user-ID  
 41015 allows a process to use *setreuid()* to swap its real and effective user IDs, but were to leave the  
 41016 saved set-user-ID unmodified, the process would then have an effective user ID equal to the  
 41017 original real user ID, and both real and saved set-user-ID would be equal to the original effective  
 41018 user ID. In that state, the real user would be unable to kill the process, even though the effective  
 41019 user ID of the process matches that of the real user, if the *kill()* behavior of `_POSIX_SAVED_IDS`  
 41020 was used. This is obviously not acceptable. The alternative choice, which is used in at least one  
 41021 implementation, is to change the saved set-user-ID to the effective user ID during most calls to  
 41022 *setreuid()*. The standard developers considered that alternative to be less correct than the  
 41023 retention of the old behavior of *kill()* in such systems. Current conforming applications shall

41024 accommodate either behavior from *kill()*, and there appears to be no strong reason for *kill()* to  
41025 check the saved set-user-ID rather than the effective user ID.

#### 41026 FUTURE DIRECTIONS

41027 None.

#### 41028 SEE ALSO

41029 *exec*, *getegid()*, *geteuid()*, *getgid()*, *getuid()*, *setegid()*, *seteuid()*, *setgid()*, *setregid()*, *setreuid()*, the  
41030 Base Definitions volume of IEEE Std. 1003.1-200x, `<sys/types.h>`, `<unistd.h>`

#### 41031 CHANGE HISTORY

41032 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 41033 Issue 4

41034 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on  
41035 XSI-conformant systems.

41036 The `<unistd.h>` header is added to the SYNOPSIS section.

41037 The following change is incorporated for alignment with the FIPS requirements:

- 41038 • All references to the saved set-user-ID are marked as extensions. This is because Issue 4  
41039 defines this mechanism as mandatory, whereas the ISO POSIX-1 standard defines that it is  
41040 only supported if `_POSIX_SAVED_IDS` is set.

#### 41041 Issue 6

41042 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

41043 The following new requirements on POSIX implementations derive from alignment with the  
41044 Single UNIX Specification:

- 41045 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
41046 required for conforming implementations of previous POSIX specifications, it was not  
41047 required for UNIX applications.
- 41048 • The functionality associated with `_POSIX_SAVED_IDS` is now mandatory. This is a FIPS  
41049 requirement.

41050 The following changes were made to align with the IEEE P1003.1a draft standard:

- 41051 • The effects of *setuid()* in processes without appropriate privileges are changed.
- 41052 • A requirement that the supplementary group list is not affected is added.

41053 **NAME**

41054 setutxent — reset user accounting database to first entry

41055 **SYNOPSIS**

41056 xSI #include <utmpx.h>

41057 void setutxent(void);

41058

41059 **DESCRIPTION**

41060 Refer to *endutxent()*.

41061 **NAME**

41062 setvbuf — assign buffering to a stream

41063 **SYNOPSIS**

41064 #include &lt;stdio.h&gt;

41065 int setvbuf(FILE \*restrict stream, char \*restrict buf, int type,  
41066 size\_t size);41067 **DESCRIPTION**41068 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
41069 conflict between the requirements described here and the ISO C standard is unintentional. This  
41070 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.41071 The *setvbuf()* function may be used after the stream pointed to by *stream* is associated with an  
41072 open file but before any other operation (other than an unsuccessful call to *setvbuf()*) is  
41073 performed on the stream. The argument *type* determines how *stream* shall be buffered, as  
41074 follows:

- 41075
- {\_IOFBF} shall cause input/output to be fully buffered.
  - {\_IOLBF} shall cause input/output to be line buffered.
  - {\_IONBF} shall cause input/output to be unbuffered.

41078 If *buf* is not a null pointer, the array it points to may be used instead of a buffer allocated by  
41079 *setvbuf()* and the argument *size* specifies the size of the array; otherwise, *size* may determine the  
41080 size of a buffer allocated by the *setvbuf()* function. The contents of the array at any time are  
41081 indeterminate.

41082 For information about streams, see Section 2.5 (on page 535).

41083 **RETURN VALUE**41084 Upon successful completion, *setvbuf()* shall return 0. Otherwise, it shall return a non-zero value  
41085 cx if an invalid value is given for *type* or if the request cannot be honored, and may set *errno* to  
41086 indicate the error.41087 **ERRORS**41088 The *setvbuf()* function may fail if:41089 cx [EBADF] The file descriptor underlying *stream* is not valid.41090 **EXAMPLES**

41091 None.

41092 **APPLICATION USAGE**41093 A common source of error is allocating buffer space as an “automatic” variable in a code block,  
41094 and then failing to close the stream in the same block.41095 With *setvbuf()*, allocating a buffer of *size* bytes does not necessarily imply that all of *size* bytes are  
41096 used for the buffer area.41097 Applications should note that many implementations only provide line buffering on input from  
41098 terminal devices.41099 **RATIONALE**

41100 None.

41101 **FUTURE DIRECTIONS**

41102 None.

41103 **SEE ALSO**41104 *fopen()*, *setbuf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**stdio.h**>41105 **CHANGE HISTORY**

41106 First released in Issue 1. Derived from Issue 1 of the SVID.

41107 **Issue 4**

41108 The second paragraph of the DESCRIPTION is now in Section 2.5 (on page 535).

41109 The [EBADF] error is marked as an extension.

41110 The APPLICATION USAGE section is expanded.

41111 The following change is incorporated for alignment with the ISO C standard:

- 41112
- This function is no longer marked as an extension.

41113 **Issue 6**

41114 Extensions beyond the ISO C standard are now marked.

41115 The *setvbuf()* prototype is updated for alignment with the ISO/IEC 9899: 1999 standard.

## 41116 NAME

41117 shm\_open — open a shared memory object (**REALTIME**)

## 41118 SYNOPSIS

41119 SHM #include &lt;sys/mman.h&gt;

41120 int shm\_open(const char \*name, int oflag, mode\_t mode);

41121

## 41122 DESCRIPTION

41123 The *shm\_open()* function establishes a connection between a shared memory object and a file  
 41124 descriptor. It creates an open file description that refers to the shared memory object and a file  
 41125 descriptor that refers to that open file description. The file descriptor is used by other functions  
 41126 to refer to that shared memory object. The *name* argument points to a string naming a shared  
 41127 memory object. It is unspecified whether the name appears in the file system and is visible to  
 41128 other functions that take path names as arguments. The *name* argument conforms to the  
 41129 construction rules for a path name. If *name* begins with the slash character, then processes calling  
 41130 *shm\_open()* with the same value of *name* refer to the same shared memory object, as long as that  
 41131 name has not been removed. If *name* does not begin with the slash character, the effect is  
 41132 implementation-defined. The interpretation of slash characters other than the leading slash  
 41133 character in *name* is implementation-defined.

41134 If successful, *shm\_open()* shall return a file descriptor for the shared memory object that is the  
 41135 lowest numbered file descriptor not currently open for that process. The open file description is  
 41136 new, and therefore the file descriptor does not share it with any other processes. It is unspecified  
 41137 whether the file offset is set. The FD\_CLOEXEC file descriptor flag associated with the new file  
 41138 descriptor is set.

41139 The file status flags and file access modes of the open file description are according to the value  
 41140 of *oflag*. The *oflag* argument is the bitwise-inclusive OR of the following flags defined in the  
 41141 header <fcntl.h>. Applications specify exactly one of the first two values (access modes) below  
 41142 in the value of *oflag*:

41143 O\_RDONLY Open for read access only.

41144 O\_RDWR Open for read or write access.

41145 Any combination of the remaining flags may be specified in the value of *oflag*:

41146 O\_CREAT If the shared memory object exists, this flag has no effect, except as noted  
 41147 under O\_EXCL below. Otherwise, the shared memory object is created; the user ID of the shared  
 41148 memory object shall be set to the effective user ID of the process; the group ID of the shared  
 41149 memory object is set to a system default group ID or to the effective group ID of the process.  
 41150 The permission bits of the shared memory object shall be set to the value of the *mode* argument  
 41151 except those set in the file mode creation mask of the process. When bits in *mode* other than  
 41152 the file permission bits are set, the effect is unspecified. The *mode* argument does not affect  
 41153 whether the shared memory object is opened for reading, for writing, or for both. The shared  
 41154 memory object has a size of zero.

41156 O\_EXCL If O\_EXCL and O\_CREAT are set, *shm\_open()* fails if the shared memory  
 41157 object exists. The check for the existence of the shared memory object and the  
 41158 creation of the object if it does not exist is atomic with respect to other  
 41159 processes executing *shm\_open()* naming the same shared memory object with  
 41160 O\_EXCL and O\_CREAT set. If O\_EXCL is set and O\_CREAT is not set, the  
 41161 result is undefined.

41162 O\_TRUNC If the shared memory object exists, and it is successfully opened O\_RDWR,  
 41163 the object shall be truncated to zero length and the mode and owner shall be  
 41164 unchanged by this function call. The result of using O\_TRUNC with  
 41165 O\_RDONLY is undefined.

41166 When a shared memory object is created, the state of the shared memory object, including all  
 41167 data associated with the shared memory object, persists until the shared memory object is  
 41168 unlinked and all other references are gone. It is unspecified whether the name and shared  
 41169 memory object state remain valid after a system reboot.

#### 41170 RETURN VALUE

41171 Upon successful completion, the *shm\_open()* function shall return a non-negative integer  
 41172 representing the lowest numbered unused file descriptor. Otherwise, it shall return `-1` and set  
 41173 *errno* to indicate the error.

#### 41174 ERRORS

41175 The *shm\_open()* function shall fail if:

41176 [EACCES] The shared memory object exists and the permissions specified by *oflag* are  
 41177 denied, or the shared memory object does not exist and permission to create  
 41178 the shared memory object is denied, or O\_TRUNC is specified and write  
 41179 permission is denied.

41180 [EEXIST] O\_CREAT and O\_EXCL are set and the named shared memory object already  
 41181 exists.

41182 [EINTR] The *shm\_open()* operation was interrupted by a signal.

41183 [EINVAL] The *shm\_open()* operation is not supported for the given name.

41184 [EMFILE] Too many file descriptors are currently in use by this process.

41185 [ENAMETOOLONG]

41186 The length of the *name* argument exceeds {PATH\_MAX} or a path name  
 41187 component is longer than {NAME\_MAX}.

41188 [ENFILE] Too many shared memory objects are currently open in the system.

41189 [ENOENT] O\_CREAT is not set and the named shared memory object does not exist.

41190 [ENOSPC] There is insufficient space for the creation of the new shared memory object.

#### 41191 EXAMPLES

41192 None.

#### 41193 APPLICATION USAGE

41194 None.

#### 41195 RATIONALE

41196 When the Memory Mapped Files option is supported, the normal *open()* call is used to obtain a  
 41197 descriptor to a file to be mapped according to existing practice with *mmap()*. When the Shared  
 41198 Memory Objects option is supported, the *shm\_open()* function is used to obtain a descriptor to  
 41199 the shared memory object to be mapped.

41200 There is ample precedent for having a file descriptor represent several types of objects. In the  
 41201 POSIX.1-1990 standard, a file descriptor can represent a file, a pipe, a FIFO, a tty, or a directory.  
 41202 Many implementations simply have an operations vector, which is indexed by the file descriptor  
 41203 type and does very different operations. Note that in some cases the file descriptor passed to  
 41204 generic operations on file descriptors are returned by *open()* or *creat()* and in some cases  
 41205 returned by alternate functions, such as *pipe()*. The latter technique is used by *shm\_open()*.

41206 Note that such shared memory objects can actually be implemented as mapped files. In both  
41207 cases, the size can be set after the open using *ftruncate()*. The *shm\_open()* function itself does not  
41208 create a shared object of a specified size because this would duplicate an extant function that set  
41209 the size of an object referenced by a file descriptor.

41210 On implementations where memory objects are implemented using the existing file system, the  
41211 *shm\_open()* function may be implemented using a macro that invokes *open()*, and the  
41212 *shm\_unlink()* function may be implemented using a macro that invokes *unlink()*.

41213 For implementations without a permanent file system, the definition of the name of the memory  
41214 objects is allowed not to survive a system reboot. Note that this allows systems with a  
41215 permanent file system to implement memory objects as data structures internal to the  
41216 implementation as well.

41217 On implementations that choose to implement memory objects using memory directly, a  
41218 *shm\_open()* followed by a *ftruncate()* and *close()* can be used to preallocate a shared memory  
41219 area and to set the size of that preallocation. This may be necessary for systems without virtual  
41220 memory hardware support in order to ensure that the memory is contiguous.

41221 The set of valid open flags to *shm\_open()* was restricted to *O\_RDONLY*, *O\_RDWR*, *O\_CREAT*,  
41222 and *O\_TRUNC* because these could be easily implemented on most memory mapping systems.  
41223 This volume of IEEE Std. 1003.1-200x is silent on the results if the implementation cannot supply  
41224 the requested file access because of implementation-defined reasons, including hardware ones.

41225 The error conditions [EACCES] and [ENOTSUP] are provided to inform the application that the  
41226 implementation cannot complete a request.

41227 [EACCES] indicates for implementation-defined reasons, probably hardware-related, that the  
41228 implementation cannot comply with a requested mode because it conflicts with another  
41229 requested mode. An example might be that an application desires to open a memory object two  
41230 times, mapping different areas with different access modes. If the implementation cannot map a  
41231 single area into a process space in two places, which would be required if different access modes  
41232 were required for the two areas, then the implementation may inform the application at the time  
41233 of the second open.

41234 [ENOTSUP] indicates for implementation-defined reasons, probably hardware-related, that the  
41235 implementation cannot comply with a requested mode at all. An example would be that the  
41236 hardware of the implementation cannot support write-only shared memory areas.

41237 On all implementations, it may be desirable to restrict the location of the memory objects to  
41238 specific file systems for performance (such as a RAM disk) or implementation-defined reasons  
41239 (shared memory supported directly only on certain file systems). The *shm\_open()* function may  
41240 be used to enforce these restrictions. There are a number of methods available to the application  
41241 to determine an appropriate name of the file or the location of an appropriate directory. One  
41242 way is from the environment via *getenv()*. Another would be from a configuration file.

41243 This volume of IEEE Std. 1003.1-200x specifies that memory objects have initial contents of zero  
41244 when created. This is consistent with current behavior for both files and newly allocated  
41245 memory. For those implementations that use physical memory, it would be possible that such  
41246 implementations could simply use available memory and give it to the process uninitialized.  
41247 This, however, is not consistent with standard behavior for the uninitialized data area, the stack,  
41248 and of course, files. Finally, it is highly desirable to set the allocated memory to zero for security  
41249 reasons. Thus, initializing memory objects to zero is required.

41250 **FUTURE DIRECTIONS**

41251 None.

41252 **SEE ALSO**

41253 *close()*, *dup()*, *exec*, *fcntl()*, *mmap()*, *shmat()*, *shmctl()*, *shmdt()*, *shm\_unlink()*, *umask()*, the Base  
41254 Definitions volume of IEEE Std. 1003.1-200x, <fcntl.h>, <sys/mman.h>

41255 **CHANGE HISTORY**

41256 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

41257 **Issue 6**41258 The *shm\_open()* function is marked as part of the Shared Memory Objects option.

41259 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
41260 implementation does not support the Shared Memory Objects option.

41261 **NAME**41262 shm\_unlink — remove a shared memory object (**REALTIME**)41263 **SYNOPSIS**41264 SHM `#include <sys/mman.h>`41265 `int shm_unlink(const char * name);`

41266

41267 **DESCRIPTION**41268 The *shm\_unlink()* function removes the name of the shared memory object named by the string  
41269 pointed to by *name*.41270 If one or more references to the shared memory object exist when the object is unlinked, the  
41271 name is removed before *shm\_unlink()* returns, but the removal of the memory object contents is  
41272 postponed until all open and map references to the shared memory object have been removed.41273 Even if the object continues to exist after the last *shm\_unlink()*, reuse of the name shall cause a  
41274 new shared memory object to be created.41275 **RETURN VALUE**41276 Upon successful completion, a value of zero shall be returned. Otherwise, a value of  $-1$  shall be  
41277 returned and *errno* set to indicate the error. If  $-1$  is returned, the named shared memory object  
41278 shall not be changed by this function call.41279 **ERRORS**41280 The *shm\_unlink()* function shall fail if:

41281 [EACCES] Permission is denied to unlink the named shared memory object.

41282 [ENAMETOOLONG]

41283 The length of the *name* argument exceeds {PATH\_MAX} or a path name  
41284 component is longer than {NAME\_MAX}.

41285 [ENOENT] The named shared memory object does not exist.

41286 **EXAMPLES**

41287 None.

41288 **APPLICATION USAGE**41289 Names of memory objects that were allocated with *open()* are deleted with *unlink()* in the usual  
41290 fashion. Names of memory objects that were allocated with *shm\_open()* are deleted with  
41291 *shm\_unlink()*. Note that the actual memory object is not destroyed until the last close and  
41292 unmap on it have occurred if it was already in use.41293 **RATIONALE**

41294 None.

41295 **FUTURE DIRECTIONS**

41296 None.

41297 **SEE ALSO**41298 *close()*, *mmap()*, *munmap()*, *shmat()*, *shmctl()*, *shmdt()*, *shm\_open()*, the Base Definitions volume  
41299 of IEEE Std. 1003.1-200x, <sys/mman.h>41300 **CHANGE HISTORY**

41301 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

41302 **Issue 6**

41303 The *shm\_unlink()* function is marked as part of the Shared Memory Objects option. |

41304 In the DESCRIPTION, text is added to clarify that reusing the same name after a *shm\_unlink()*  
41305 will not attach to the old shared memory object.

41306 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
41307 implementation does not support the Shared Memory Objects option. |

41308 **NAME**

41309 shmat — XSI shared memory attach operation

41310 **SYNOPSIS**41311 XSI 

```
#include <sys/shm.h>
```

41312 

```
void *shmat(int shmid, const void *shmaddr, int shmflg);
```

41313

41314 **DESCRIPTION**

41315 The *shmat()* function operates on XSI shared memory (see the Base Definitions volume of  
 41316 IEEE Std. 1003.1-200x, Section 3.342, Shared Memory Object). It is unspecified whether this  
 41317 function interoperates with the realtime interprocess communication facilities defined in Section  
 41318 2.8 (on page 543).

41319 The *shmat()* function attaches the shared memory segment associated with the shared memory  
 41320 identifier specified by *shmid* to the address space of the calling process. The segment is attached  
 41321 at the address specified by one of the following criteria:

- 41322 • If *shmaddr* is a null pointer, the segment is attached at the first available address as selected  
 41323 by the system.
- 41324 • If *shmaddr* is not a null pointer and (*shmflg* &SHM\_RND) is non-zero, the segment is attached  
 41325 at the address given by (*shmaddr* - ((*uintptr\_t*)*shmaddr* %SHMLBA)). The character '%' is the  
 41326 C-language remainder operator.
- 41327 • If *shmaddr* is not a null pointer and (*shmflg* &SHM\_RND) is 0, the segment is attached at the  
 41328 address given by *shmaddr*.
- 41329 • The segment is attached for reading if (*shmflg* &SHM\_RDONLY) is non-zero and the calling  
 41330 process has read permission; otherwise, if it is 0 and the calling process has read and write  
 41331 permission, the segment is attached for reading and writing.

41332 **RETURN VALUE**

41333 Upon successful completion, *shmat()* shall increment the value of *shm\_nattach* in the data  
 41334 structure associated with the shared memory ID of the attached shared memory segment and  
 41335 return the segment's start address.

41336 Otherwise, the shared memory segment shall not be attached, *shmat()* shall return -1, and *errno*  
 41337 shall be set to indicate the error.

41338 **ERRORS**41339 The *shmat()* function shall fail if:

- |                                           |          |                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 41340<br>41341                            | [EACCES] | Operation permission is denied to the calling process; see Section 2.7 (on page 541).                                                                                                                                                                                                                                                                                                                              |
| 41342<br>41343<br>41344<br>41345<br>41346 | [EINVAL] | The value of <i>shmid</i> is not a valid shared memory identifier, the <i>shmaddr</i> is not a null pointer, and the value of ( <i>shmaddr</i> - (( <i>uintptr_t</i> ) <i>shmaddr</i> %SHMLBA)) is an illegal address for attaching shared memory; or the <i>shmaddr</i> is not a null pointer, ( <i>shmflg</i> &SHM_RND) is 0, and the value of <i>shmaddr</i> is an illegal address for attaching shared memory. |
| 41347<br>41348                            | [EMFILE] | The number of shared memory segments attached to the calling process would exceed the system-imposed limit.                                                                                                                                                                                                                                                                                                        |
| 41349<br>41350                            | [ENOMEM] | The available data space is not large enough to accommodate the shared memory segment.                                                                                                                                                                                                                                                                                                                             |

41351 **EXAMPLES**

41352 None.

41353 **APPLICATION USAGE**

41354 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
41355 Application developers who need to use IPC should design their applications so that modules  
41356 using the IPC routines described in Section 2.7 (on page 541) can be easily modified to use the  
41357 alternative interfaces.

41358 **RATIONALE**

41359 None.

41360 **FUTURE DIRECTIONS**

41361 None.

41362 **SEE ALSO**

41363 *exec*, *exit()*, *fork()*, *shmctl()*, *shmdt()*, *shmget()*, *shm\_open()*, *shm\_unlink()*, the Base Definitions  
41364 volume of IEEE Std. 1003.1-200x, <sys/shm.h>, Section 2.7 (on page 541)

41365 **CHANGE HISTORY**

41366 First released in Issue 2. Derived from Issue 2 of the SVID.

41367 **Issue 4**

41368 The function is no longer marked as OPTIONAL FUNCTIONALITY.

41369 The &lt;sys/types.h&gt; and &lt;sys/ipc.h&gt; headers are removed from the SYNOPSIS section.

41370 The type of argument *shmaddr* is changed from **char\*** to **const void\***.

41371 The [ENOSYS] error is removed from the ERRORS section.

41372 The DESCRIPTION is clarified in several places.

41373 A FUTURE DIRECTIONS section is added warning application developers about migration to  
41374 IEEE 1003.4 interfaces for interprocess communication.

41375 **Issue 5**

41376 Moved from SHARED MEMORY to BASE.

41377 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
41378 DIRECTIONS to a new APPLICATION USAGE section.

41379 **Issue 6**

41380 The Open Group corrigenda item U021/13 has been applied.

## 41381 NAME

41382 shmctl — XSI shared memory control operations

## 41383 SYNOPSIS

41384 XSI 

```
#include <sys/shm.h>
```

41385 

```
int shmctl(int shmid, int cmd, struct shm_id_ds *buf);
```

41386

## 41387 DESCRIPTION

41388 The *shmctl()* function operates on XSI shared memory (see the Base Definitions volume of  
 41389 IEEE Std. 1003.1-200x, Section 3.342, Shared Memory Object). It is unspecified whether this  
 41390 function interoperates with the realtime interprocess communication facilities defined in Section  
 41391 2.8 (on page 543).

41392 The *shmctl()* function provides a variety of shared memory control operations as specified by  
 41393 *cmd*. The following values for *cmd* are available:

41394 IPC\_STAT Place the current value of each member of the **shm\_id\_ds** data structure  
 41395 associated with *shmid* into the structure pointed to by *buf*. The contents of the  
 41396 structure are defined in `<sys/shm.h>`.

41397 IPC\_SET Set the value of the following members of the **shm\_id\_ds** data structure  
 41398 associated with *shmid* to the corresponding value found in the structure  
 41399 pointed to by *buf*:

41400 shm\_perm.uid  
 41401 shm\_perm.gid  
 41402 shm\_perm.mode Low-order nine bits.

41403 IPC\_SET can only be executed by a process that has an effective user ID equal  
 41404 to either that of a process with appropriate privileges or to the value of  
 41405 *shm\_perm.cuid* or *shm\_perm.uid* in the **shm\_id\_ds** data structure associated with  
 41406 *shmid*.

41407 IPC\_RMID Remove the shared memory identifier specified by *shmid* from the system and  
 41408 destroy the shared memory segment and **shm\_id\_ds** data structure associated  
 41409 with it. IPC\_RMID can only be executed by a process that has an effective user  
 41410 ID equal to either that of a process with appropriate privileges or to the value  
 41411 of *shm\_perm.cuid* or *shm\_perm.uid* in the **shm\_id\_ds** data structure associated  
 41412 with *shmid*.

## 41413 RETURN VALUE

41414 Upon successful completion, *shmctl()* shall return 0; otherwise, it shall return -1 and set *errno* to  
 41415 indicate the error.

## 41416 ERRORS

41417 The *shmctl()* function shall fail if:

41418 [EACCES] The argument *cmd* is equal to IPC\_STAT and the calling process does not have  
 41419 read permission; see Section 2.7 (on page 541).

41420 [EINVAL] The value of *shmid* is not a valid shared memory identifier, or the value of *cmd*  
 41421 is not a valid command.

41422 [EPERM] The argument *cmd* is equal to IPC\_RMID or IPC\_SET and the effective user ID  
 41423 of the calling process is not equal to that of a process with appropriate  
 41424 privileges and it is not equal to the value of *shm\_perm.cuid* or *shm\_perm.uid* in  
 41425 the data structure associated with *shmid*.

41426 The *shmctl()* function may fail if:

41427 [EOVERFLOW] The *cmd* argument is IPC\_STAT and the *gid* or *uid* value is too large to be  
41428 stored in the structure pointed to by the *buf* argument.

41429 **EXAMPLES**

41430 None.

41431 **APPLICATION USAGE**

41432 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
41433 Application developers who need to use IPC should design their applications so that modules  
41434 using the IPC routines described in Section 2.7 (on page 541) can be easily modified to use the  
41435 alternative interfaces.

41436 **RATIONALE**

41437 None.

41438 **FUTURE DIRECTIONS**

41439 None.

41440 **SEE ALSO**

41441 *shmat()*, *shmdt()*, *shmget()*, *shm\_open()*, *shm\_unlink()*, the Base Definitions volume of  
41442 IEEE Std. 1003.1-200x, <sys/shm.h>, Section 2.7 (on page 541)

41443 **CHANGE HISTORY**

41444 First released in Issue 2. Derived from Issue 2 of the SVID.

41445 **Issue 4**

41446 The function is no longer marked as OPTIONAL FUNCTIONALITY.

41447 The <sys/types.h> and <sys/ipc.h> headers are removed from the SYNOPSIS section.

41448 The [ENOSYS] error is removed from the ERRORS section.

41449 A FUTURE DIRECTIONS section is added warning application developers about migration to  
41450 IEEE 1003.4 interfaces for interprocess communication.

41451 **Issue 4, Version 2**

41452 The ERRORS section is updated for X/OPEN UNIX conformance to include [EOVERFLOW] as  
41453 an optional error.

41454 **Issue 5**

41455 Moved from SHARED MEMORY to BASE.

41456 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
41457 DIRECTIONS to a new APPLICATION USAGE section.

41458 **NAME**

41459 shmdt — XSI shared memory detach operation

41460 **SYNOPSIS**41461 XSI 

```
#include <sys/shm.h>
```

41462 

```
int shmdt(const void *shmaddr);
```

41463

41464 **DESCRIPTION**

41465 The *shmdt()* function operates on XSI shared memory (see the Base Definitions volume of  
41466 IEEE Std. 1003.1-200x, Section 3.342, Shared Memory Object). It is unspecified whether this  
41467 function interoperates with the realtime interprocess communication facilities defined in Section  
41468 2.8 (on page 543).

41469 The *shmdt()* function detaches the shared memory segment located at the address specified by  
41470 *shmaddr* from the address space of the calling process.

41471 **RETURN VALUE**

41472 Upon successful completion, *shmdt()* shall decrement the value of *shm\_nattch* in the data  
41473 structure associated with the shared memory ID of the attached shared memory segment and  
41474 return 0.

41475 Otherwise, the shared memory segment shall not be detached, *shmdt()* shall return  $-1$ , and *errno*  
41476 shall be set to indicate the error.

41477 **ERRORS**41478 The *shmdt()* function shall fail if:

41479 [EINVAL] The value of *shmaddr* is not the data segment start address of a shared  
41480 memory segment.

41481 **EXAMPLES**

41482 None.

41483 **APPLICATION USAGE**

41484 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
41485 Application developers who need to use IPC should design their applications so that modules  
41486 using the IPC routines described in Section 2.7 (on page 541) can be easily modified to use the  
41487 alternative interfaces.

41488 **RATIONALE**

41489 None.

41490 **FUTURE DIRECTIONS**

41491 None.

41492 **SEE ALSO**

41493 *exec*, *exit()*, *fork()*, *shmat()*, *shmctl()*, *shmget()*, *shm\_open()*, *shm\_unlink()*, the Base Definitions  
41494 volume of IEEE Std. 1003.1-200x, <sys/shm.h>, Section 2.7 (on page 541)

41495 **CHANGE HISTORY**

41496 First released in Issue 2. Derived from Issue 2 of the SVID.

41497 **Issue 4**

41498 The function is no longer marked as OPTIONAL FUNCTIONALITY.

41499 The &lt;sys/types.h&gt; and &lt;sys/ipc.h&gt; headers are removed from the SYNOPSIS section.

41500 The type of argument *shmaddr* is changed from **char\*** to **const void\***.

- 41501 The DESCRIPTION is clarified in several places.
- 41502 The [ENOSYS] error is removed from the ERRORS section.
- 41503 A FUTURE DIRECTIONS section is added warning application developers about migration to
- 41504 IEEE 1003.4 interfaces for interprocess communication.
- 41505 **Issue 5**
- 41506 Moved from SHARED MEMORY to BASE.
- 41507 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE
- 41508 DIRECTIONS to a new APPLICATION USAGE section.

## 41509 NAME

41510 shmget — get XSI shared memory segment

## 41511 SYNOPSIS

41512 XSI #include &lt;sys/shm.h&gt;

41513 int shmget(key\_t key, size\_t size, int shmflg);

41514

## 41515 DESCRIPTION

41516 The *shmget()* function operates on XSI shared memory (see the Base Definitions volume of  
 41517 IEEE Std. 1003.1-200x, Section 3.342, Shared Memory Object). It is unspecified whether this  
 41518 function interoperates with the realtime interprocess communication facilities defined in Section  
 41519 2.8 (on page 543).

41520 The *shmget()* function shall return the shared memory identifier associated with *key*.

41521 A shared memory identifier, associated data structure, and shared memory segment of at least  
 41522 *size* bytes (see <sys/shm.h>) are created for *key* if one of the following is true:

- 41523 • The argument *key* is equal to `IPC_PRIVATE`.
- 41524 • The argument *key* does not already have a shared memory identifier associated with it and  
 41525 (*shmflg* & `IPC_CREAT`) is non-zero.

41526 Upon creation, the data structure associated with the new shared memory identifier shall be  
 41527 initialized as follows:

- 41528 • The values of *shm\_perm.cuid*, *shm\_perm.uid*, *shm\_perm.cgid*, and *shm\_perm.gid* are set equal to  
 41529 the effective user ID and effective group ID, respectively, of the calling process.
- 41530 • The low-order nine bits of *shm\_perm.mode* are set equal to the low-order nine bits of *shmflg*.  
 41531 The value of *shm\_segsz* is set equal to the value of *size*.
- 41532 • The values of *shm\_lpid*, *shm\_nattch*, *shm\_atime*, and *shm\_dtime* are set equal to 0.
- 41533 • The value of *shm\_ctime* is set equal to the current time.

41534 When the shared memory segment is created, it shall be initialized with all zero values.

## 41535 RETURN VALUE

41536 Upon successful completion, *shmget()* shall return a non-negative integer, namely a shared  
 41537 memory identifier; otherwise, it shall return `-1` and set *errno* to indicate the error.

## 41538 ERRORS

41539 The *shmget()* function shall fail if:

- |                                  |          |                                                                                                                                                                                                                                                                                 |
|----------------------------------|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 41540<br>41541<br>41542          | [EACCES] | A shared memory identifier exists for <i>key</i> but operation permission as specified by the low-order nine bits of <i>shmflg</i> would not be granted; see Section 2.7 (on page 541).                                                                                         |
| 41543<br>41544                   | [EEXIST] | A shared memory identifier exists for the argument <i>key</i> but ( <i>shmflg</i> & <code>IPC_CREAT</code> ) && ( <i>shmflg</i> & <code>IPC_EXCL</code> ) is non-zero.                                                                                                          |
| 41545<br>41546<br>41547<br>41548 | [EINVAL] | The value of <i>size</i> is less than the system-imposed minimum or greater than the system-imposed maximum, or a shared memory identifier exists for the argument <i>key</i> but the size of the segment associated with it is less than <i>size</i> and <i>size</i> is not 0. |
| 41549<br>41550                   | [ENOENT] | A shared memory identifier does not exist for the argument <i>key</i> and ( <i>shmflg</i> & <code>IPC_CREAT</code> ) is 0.                                                                                                                                                      |

- 41551 [ENOMEM] A shared memory identifier and associated shared memory segment shall be  
 41552 created, but the amount of available physical memory is not sufficient to fill  
 41553 the request.
- 41554 [ENOSPC] A shared memory identifier is to be created, but the system-imposed limit on  
 41555 the maximum number of allowed shared memory identifiers system-wide  
 41556 would be exceeded.
- 41557 **EXAMPLES**
- 41558 None.
- 41559 **APPLICATION USAGE**
- 41560 The POSIX Realtime Extension defines alternative interfaces for interprocess communication.  
 41561 Application developers who need to use IPC should design their applications so that modules  
 41562 using the IPC routines described in Section 2.7 (on page 541) can be easily modified to use the  
 41563 alternative interfaces.
- 41564 **RATIONALE**
- 41565 None.
- 41566 **FUTURE DIRECTIONS**
- 41567 None.
- 41568 **SEE ALSO**
- 41569 *shmat()*, *shmctl()*, *shmdt()*, *shm\_open()*, *shm\_unlink()*, the Base Definitions volume of  
 41570 IEEE Std. 1003.1-200x, <sys/shm.h>, Section 2.7 (on page 541)
- 41571 **CHANGE HISTORY**
- 41572 First released in Issue 2. Derived from Issue 2 of the SVID.
- 41573 **Issue 4**
- 41574 The function is no longer marked as OPTIONAL FUNCTIONALITY.
- 41575 The <sys/types.h> and <sys/ipc.h> headers are removed from the SYNOPSIS section.
- 41576 The [ENOSYS] error is removed from the ERRORS section.
- 41577 A FUTURE DIRECTIONS section is added warning application developers about migration to  
 41578 IEEE 1003.4 interfaces for interprocess communication.
- 41579 **Issue 5**
- 41580 Moved from SHARED MEMORY to BASE.
- 41581 The note about use of POSIX Realtime Extension IPC routines has been moved from FUTURE  
 41582 DIRECTIONS to a new APPLICATION USAGE section.

41583 **NAME**

41584 shutdown — shut down socket send and receive operations

41585 **SYNOPSIS**

41586 #include &lt;sys/socket.h&gt;

41587 int shutdown(int *socket*, int *how*);41588 **DESCRIPTION**41589 The *shutdown()* function shall cause all or part of a full-duplex connection on the socket  
41590 associated with the file descriptor *socket* to be shut down.41591 The *shutdown()* function takes the following arguments:

|       |               |                                                            |
|-------|---------------|------------------------------------------------------------|
| 41592 | <i>socket</i> | Specifies the file descriptor of the socket.               |
| 41593 | <i>how</i>    | Specifies the type of shutdown. The values are as follows: |
| 41594 | SHUT_RD       | Disables further receive operations.                       |
| 41595 | SHUT_WR       | Disables further send operations.                          |
| 41596 | SHUT_RDWR     | Disables further send and receive operations.              |

41597 The *shutdown()* function disables subsequent send and/or receive operations on a socket,  
41598 depending on the value of the *how* argument.41599 **RETURN VALUE**41600 Upon successful completion, *shutdown()* shall return 0; otherwise, -1 shall be returned and *errno*  
41601 set to indicate the error.41602 **ERRORS**41603 The *shutdown()* function shall fail if:

|       |            |                                                            |
|-------|------------|------------------------------------------------------------|
| 41604 | [EBADF]    | The <i>socket</i> argument is not a valid file descriptor. |
| 41605 | [EINVAL]   | The <i>how</i> argument is invalid.                        |
| 41606 | [ENOTCONN] | The socket is not connected.                               |
| 41607 | [ENOTSOCK] | The <i>socket</i> argument does not refer to a socket.     |

41608 The *shutdown()* function may fail if:

41609 [ENOBUFS] Insufficient resources were available in the system to perform the operation. |

41610 **EXAMPLES**

41611 None.

41612 **APPLICATION USAGE**

41613 None.

41614 **RATIONALE**

41615 None.

41616 **FUTURE DIRECTIONS**

41617 None.

41618 **SEE ALSO**41619 *getsockopt()*, *read()*, *recv()*, *recvfrom()*, *recvmsg()*, *select()*, *send()*, *sendto()*, *setsockopt()*, *socket()*,  
41620 *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/socket.h> |

41621 **CHANGE HISTORY**

41622 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

## 41623 NAME

41624 sigaction — examine and change signal action

## 41625 SYNOPSIS

41626 #include &lt;signal.h&gt;

41627 int sigaction(int *sig*, const struct sigaction \*restrict *act*,  
41628 struct sigaction \*restrict *oact*);

## 41629 DESCRIPTION

41630 The *sigaction()* function allows the calling process to examine and/or specify the action to be  
41631 associated with a specific signal. The argument *sig* specifies the signal; acceptable values are  
41632 defined in <signal.h>.41633 The structure **sigaction**, used to describe an action to be taken, is defined in the header  
41634 <signal.h> to include at least the following members:

41635

41636

| Member Type                           | Member Name                         | Description                                                                                                                                |
|---------------------------------------|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| void(*) (int)<br>sigset_t             | <i>sa_handler</i><br><i>sa_mask</i> | SIG_DFL, SIG_IGN, or pointer to a function.<br>Additional set of signals to be blocked<br>during execution of signal-catching<br>function. |
| int                                   | <i>sa_flags</i>                     | Special flags to affect behavior of signal.                                                                                                |
| void(*) (int,<br>siginfo_t *, void *) | <i>sa_sigaction</i>                 | Signal-catching function.                                                                                                                  |

41643

41644 If the argument *act* is not a null pointer, it points to a structure specifying the action to be  
41645 associated with the specified signal. If the argument *oact* is not a null pointer, the action  
41646 previously associated with the signal is stored in the location pointed to by the argument *oact*. If  
41647 the argument *act* is a null pointer, signal handling is unchanged; thus, the call can be used to  
41648 enquire about the current handling of a given signal. The SIGKILL and SIGSTOP signals shall  
41649 not be added to the signal mask using this mechanism; this restriction shall be enforced by the  
41650 system without causing an error to be indicated.41651 If the SA\_SIGINFO flag (see below) is cleared in the *sa\_flags* field of the **sigaction** structure, the  
41652 *sa\_handler* field identifies the action to be associated with the specified signal. If the  
41653 SA\_SIGINFO flag is set in the *sa\_flags* field, and the implementation supports the Realtime  
41654 Signals Extension option or the X/Open System Interfaces Extension option, the *sa\_sigaction*  
41655 field specifies a signal-catching function. If the SA\_SIGINFO bit is cleared and the *sa\_handler*  
41656 field specifies a signal-catching function, or if the SA\_SIGINFO bit is set, the *sa\_mask* field  
41657 identifies a set of signals that shall be added to the signal mask of the thread before the signal-  
41658 catching function is invoked. If the *sa\_handler* field specifies a signal-catching function, the  
41659 *sa\_mask* field identifies a set of signals that shall be added to the process' signal mask before the  
41660 signal-catching function is invoked.41661 The *sa\_flags* field can be used to modify the behavior of the specified signal.41662 The following flags, defined in the header <signal.h>, can be set in *sa\_flags*:

41663 SA\_NOCLDSTOP Do not generate SIGCHLD when children stop.

41664 If *sig* is SIGCHLD and the SA\_NOCLDSTOP flag is not set in *sa\_flags*, and  
41665 the implementation supports the SIGCHLD signal, then a SIGCHLD  
41666 signal shall be generated for the calling process whenever any of its child  
41667 processes stop. If *sig* is SIGCHLD and the SA\_NOCLDSTOP flag is set in  
41668 *sa\_flags*, then the implementation shall not generate a SIGCHLD signal in

|               |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 41669         |              | this way.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 41670 XSI     | SA_ONSTACK   | If set and an alternate signal stack has been declared with <i>sigaltstack()</i> or <i>sigstack()</i> , the signal shall be delivered to the calling process on that stack. Otherwise, the signal shall be delivered on the current stack.                                                                                                                                                                                                                                                                                                                                |
| 41671         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41672         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41673 XSI     | SA_RESETHAND | If set, the disposition of the signal shall be reset to SIG_DFL and the SA_SIGINFO flag shall be cleared on entry to the signal handler.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| 41674         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41675         |              | <b>Note:</b> SIGILL and SIGTRAP cannot be automatically reset when delivered; the system silently enforces this restriction.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 41676         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41677         |              | Otherwise, the disposition of the signal shall not be modified on entry to the signal handler.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 41678         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41679         |              | In addition, if this flag is set, <i>sigaction()</i> behaves as if the SA_NODEFER flag were also set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 41680         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41681 XSI     | SA_RESTART   | This flag affects the behavior of interruptible functions; that is, those specified to fail with <i>errno</i> set to [EINTR]. If set, and a function specified as interruptible is interrupted by this signal, the function shall restart and shall not fail with [EINTR] unless otherwise specified. If the flag is not set, interruptible functions interrupted by this signal shall fail with <i>errno</i> set to [EINTR].                                                                                                                                             |
| 41682         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41683         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41684         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41685         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41686         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41687         | SA_SIGINFO   | If cleared and the signal is caught, the signal-catching function shall be entered as:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 41688         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41689         |              | <pre>void func(int signo);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 41690         |              | where <i>signo</i> is the only argument to the signal catching function. In this case, the application shall use the <i>sa_handler</i> member to describe the signal catching function and the application shall not modify the <i>sa_sigaction</i> member.                                                                                                                                                                                                                                                                                                               |
| 41691         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41692         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41693         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41694 XSI RTS |              | If SA_SIGINFO is set and the signal is caught, the signal-catching function shall be entered as:                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 41695         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41696         |              | <pre>void func(int signo, siginfo_t *info, void *context);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 41697         |              | where two additional arguments are passed to the signal catching function. The second argument shall point to an object of type <b>siginfo_t</b> explaining the reason why the signal was generated; the third argument can be cast to a pointer to an object of type <b>ucontext_t</b> to refer to the receiving process' context that was interrupted when the signal was delivered. In this case, the application shall use the <i>sa_sigaction</i> member to describe the signal catching function and the application shall not modify the <i>sa_handler</i> member. |
| 41698         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41699         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41700         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41701         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41702         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41703         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41704         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41705         |              | The <i>si_signo</i> member contains the system-generated signal number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 41706 XSI     |              | The <i>si_errno</i> member may contain implementation-defined additional error information; if non-zero, it contains an error number identifying the condition that caused the signal to be generated.                                                                                                                                                                                                                                                                                                                                                                    |
| 41707         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41708         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41709         |              | The <i>si_code</i> member contains a code identifying the cause of the signal. If the value of <i>si_code</i> is less than or equal to 0, then the signal was generated by a process and <i>si_pid</i> and <i>si_uid</i> , respectively, indicate the process ID and the real user ID of the sender. The <b>&lt;signal.h&gt;</b> header description contains information about the signal specific contents of the                                                                                                                                                        |
| 41710         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41711         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41712         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| 41713         |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

41714 elements of the **siginfo\_t** type.

41715 XSI **SA\_NOCLDWAIT** If set, and *sig* equals SIGCHLD, child processes of the calling processes  
41716 shall not be transformed into zombie processes when they terminate. If  
41717 the calling process subsequently waits for its children, and the process  
41718 has no unwaited-for children that were transformed into zombie  
41719 processes, it shall block until all of its children terminate, and *wait()*,  
41720 *waitid()*, and *waitpid()* shall fail and set *errno* to [ECHILD]. Otherwise,  
41721 terminating child processes shall be transformed into zombie processes,  
41722 unless SIGCHLD is set to SIG\_IGN.

41723 XSI **SA\_NODEFER** If set and *sig* is caught, *sig* shall not be added to the process' signal mask  
41724 on entry to the signal handler unless it is included in *sa\_mask*. Otherwise,  
41725 *sig* shall always be added to the process' signal mask on entry to the  
41726 signal handler.

41727 When a signal is caught by a signal-catching function installed by *sigaction()*, a new signal mask  
41728 is calculated and installed for the duration of the signal-catching function (or until a call to either  
41729 *sigprocmask()* or *sigsuspend()* is made). This mask is formed by taking the union of the current  
41730 XSI signal mask and the value of the *sa\_mask* for the signal being delivered unless SA\_NODEFER or  
41731 SA\_RESETHAND is set, and then including the signal being delivered. If and when the user's  
41732 signal handler returns normally, the original signal mask is restored.

41733 Once an action is installed for a specific signal, it remains installed until another action is  
41734 XSI explicitly requested (by another call to *sigaction()*), until the SA\_RESETHAND flag causes  
41735 resetting of the handler, or until one of the *exec* functions is called.

41736 If the previous action for *sig* had been established by *signal()*, the values of the fields returned in  
41737 the structure pointed to by *oact* are unspecified, and in particular *oact->sa\_handler* is not  
41738 necessarily the same value passed to *signal()*. However, if a pointer to the same structure or a  
41739 copy thereof is passed to a subsequent call to *sigaction()* via the *act* argument, handling of the  
41740 signal shall be as if the original call to *signal()* were repeated.

41741 If *sigaction()* fails, no new signal handler is installed.

41742 It is unspecified whether an attempt to set the action for a signal that cannot be caught or  
41743 ignored to SIG\_DFL is ignored or causes an error to be returned with *errno* set to [EINVAL].

41744 If SA\_SIGINFO is not set in *sa\_flags*, then the disposition of subsequent occurrences of *sig* when  
41745 it is already pending is implementation-defined; the signal-catching function shall be invoked  
41746 RTS with a single argument. If the implementation supports the Realtime Signals Extension option,  
41747 and if SA\_SIGINFO is set in *sa\_flags*, then subsequent occurrences of *sig* generated by *sigqueue()*  
41748 or as a result of any signal-generating function that supports the specification of an application-  
41749 defined value (when *sig* is already pending) shall be queued in FIFO order until delivered or  
41750 accepted; the signal-catching function shall be invoked with three arguments. The application  
41751 specified value is passed to the signal-catching function as the *si\_value* member of the **siginfo\_t**  
41752 structure.

41753 The result of the use of *sigaction()* and a *sigwait()* function concurrently within a process on the  
41754 same signal is unspecified.

41755 **RETURN VALUE**

41756 Upon successful completion, *sigaction()* shall return 0; otherwise, -1 shall be returned, *errno* shall  
41757 be set to indicate the error, and no new signal-catching function shall be installed.

41758 **ERRORS**41759 The *sigaction()* function shall fail if:41760 [EINVAL] The *sig* argument is not a valid signal number or an attempt is made to catch a  
41761 signal that cannot be caught or ignore a signal that cannot be ignored.41762 [ENOTSUP] The SA\_SIGINFO bit flag is set in the *sa\_flags* field of the **sigaction** structure,  
41763 and the implementation does not support either the Realtime Signals  
41764 Extension option, or the X/Open System Interfaces Extension option.41765 The *sigaction()* function may fail if:41766 [EINVAL] An attempt was made to set the action to SIG\_DFL for a signal that cannot be  
41767 caught or ignored (or both).41768 **EXAMPLES**

41769 None.

41770 **APPLICATION USAGE**41771 The *sigaction()* function supersedes the *signal()* function, and should be used in preference. In  
41772 particular, *sigaction()* and *signal()* should not be used in the same process to control the same  
41773 signal. The behavior of reentrant functions, as defined in the DESCRIPTION, is as specified by  
41774 this volume of IEEE Std. 1003.1-200x, regardless of invocation from a signal-catching function.  
41775 This is the only intended meaning of the statement that reentrant functions may be used in  
41776 signal-catching functions without restrictions. Applications must still consider all effects of such  
41777 functions on such things as data structures, files, and process state. In particular, application  
41778 writers need to consider the restrictions on interactions when interrupting *sleep()* and  
41779 interactions among multiple handles for a file description. The fact that any specific function is  
41780 listed as reentrant does not necessarily mean that invocation of that function from a signal-  
41781 catching function is recommended.41782 In order to prevent errors arising from interrupting non-reentrant function calls, applications  
41783 should protect calls to these functions either by blocking the appropriate signals or through the  
41784 use of some programmatic semaphore (see *semget()*, *sem\_init()*, *sem\_open()*, and so on). Note in  
41785 particular that even the “safe” functions may modify *errno*; the signal-catching function, if not  
41786 executing as an independent thread, may want to save and restore its value. Naturally, the same  
41787 principles apply to the reentrancy of application routines and asynchronous data access. Note  
41788 that *longjmp()* and *siglongjmp()* are not in the list of reentrant functions. This is because the code  
41789 executing after *longjmp()* and *siglongjmp()* can call any unsafe functions with the same danger as  
41790 calling those unsafe functions directly from the signal handler. Applications that use *longjmp()*  
41791 and *siglongjmp()* from within signal handlers require rigorous protection in order to be portable.  
41792 Many of the other functions that are excluded from the list are traditionally implemented using  
41793 either *malloc()* or *free()* functions or the standard I/O library, both of which traditionally use  
41794 data structures in a non-reentrant manner. Because any combination of different functions using  
41795 a common data structure can cause reentrancy problems, this volume of IEEE Std. 1003.1-200x  
41796 does not define the behavior when any unsafe function is called in a signal handler that  
41797 interrupts an unsafe function.41798 If the signal occurs other than as the result of calling *abort()*, *kill()*, or *raise()*, the behavior is  
41799 undefined if the signal handler calls any function in the standard library other than one of the  
41800 functions listed in the table above or refers to any object with static storage duration other than  
41801 by assigning a value to a static storage duration variable of type **volatile sig\_atomic\_t**.  
41802 Furthermore, if such a call fails, the value of *errno* is indeterminate.41803 Usually, the signal is executed on the stack that was in effect before the signal was delivered. An  
41804 alternate stack may be specified to receive a subset of the signals being caught.

41805 When the signal handler returns, the receiving process resumes execution at the point it was  
41806 interrupted unless the signal handler makes other arrangements. If *longjmp()* or *\_longjmp()* is  
41807 used to leave the signal handler, then the signal mask must be explicitly restored by the process.

41808 This volume of IEEE Std. 1003.1-200x defines the third argument of a signal handling function  
41809 when SA\_SIGINFO is set as a **void\*** instead of a **ucontext\_t\***, but without requiring type  
41810 checking. New applications should explicitly cast the third argument of the signal handling  
41811 function to **ucontext\_t\***.

41812 The BSD optional four argument signal handling function is not supported by this volume of  
41813 IEEE Std. 1003.1-200x. The BSD declaration would be:

```
41814 void handler(int sig, int code, struct sigcontext *scp,
41815 char *addr);
```

41816 where *sig* is the signal number, *code* is additional information on certain signals, *scp* is a pointer  
41817 to the sigcontext structure, and *addr* is additional address information. Much the same  
41818 information is available in the objects pointed to by the second argument of the signal handler  
41819 specified when SA\_SIGINFO is set.

#### 41820 RATIONALE

41821 Although this volume of IEEE Std. 1003.1-200x requires that signals that cannot be ignored shall  
41822 not be added to the signal mask when a signal-catching function is entered, there is no explicit  
41823 requirement that subsequent calls to *sigaction()* reflect this in the information returned in the *oact*  
41824 argument. In other words, if SIGKILL is included in the *sa\_mask* field of *act*, it is unspecified  
41825 whether or not a subsequent call to *sigaction()* returns with SIGKILL included in the *sa\_mask*  
41826 field of *oact*.

41827 The SA\_NOCLDSTOP flag, when supplied in the *act->sa\_flags* parameter, allows overloading  
41828 SIGCHLD with the System V semantics that each SIGCLD signal indicates a single terminated  
41829 child. Most portable applications that catch SIGCHLD are expected to install signal-catching  
41830 functions that repeatedly call the *waitpid()* function with the WNOHANG flag set, acting on  
41831 each child for which status is returned, until *waitpid()* returns zero. If stopped children are not of  
41832 interest, the use of the SA\_NOCLDSTOP flag can prevent the overhead from invoking the  
41833 signal-catching routine when they stop.

41834 Some historical implementations also define other mechanisms for stopping processes, such as  
41835 the *ptrace()* function. These implementations usually do not generate a SIGCHLD signal when  
41836 processes stop due to this mechanism; however, that is beyond the scope of this volume of  
41837 IEEE Std. 1003.1-200x.

41838 This volume of IEEE Std. 1003.1-200x requires that calls to *sigaction()* that supply a NULL *act*  
41839 argument succeed, even in the case of signals that cannot be caught or ignored (that is, SIGKILL  
41840 or SIGSTOP). The System V *signal()* and BSD *sigvec()* functions return [EINVAL] in these cases  
41841 and, in this respect, their behavior varies from *sigaction()*.

41842 This volume of IEEE Std. 1003.1-200x requires that *sigaction()* properly save and restore a signal  
41843 action set up by the ISO C standard *signal()* function. However, there is no guarantee that the  
41844 reverse is true, nor could there be given the greater amount of information conveyed by the  
41845 **sigaction** structure. Because of this, applications should avoid using both functions for the same  
41846 signal in the same process. Since this cannot always be avoided in case of general-purpose  
41847 library routines, they should always be implemented with *sigaction()*.

41848 It was intended that the *signal()* function should be implementable as a library routine using  
41849 *sigaction()*.

41850 The POSIX Realtime Extension extends the *sigaction()* function as specified by the POSIX.1-1990  
41851 standard to allow the application to request on a per-signal basis via an additional signal action

41852 flag that the extra parameters, including the application-defined signal value, if any, be passed  
41853 to the signal-catching function.

#### 41854 FUTURE DIRECTIONS

41855 The *fpathconf()* function is marked as an extension in the list of safe functions because it is not  
41856 included in the corresponding list in the ISO POSIX-1 standard, but it is expected to be added in  
41857 a future version.

#### 41858 SEE ALSO

41859 Section 2.4 (on page 528), *bsd\_signal()*, *kill()*, *\_longjmp()*, *longjmp()*, *raise()*, *semget()*, *sem\_init()*,  
41860 *sem\_open()*, *sigaddset()*, *sigaltstack()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*, *signal()*,  
41861 *sigprocmask()*, *sigsuspend()*, *wait()*, *waitid()*, *waitpid()*, the Base Definitions volume of  
41862 IEEE Std. 1003.1-200x, <**signal.h**>, <**ucontext.h**>

#### 41863 CHANGE HISTORY

41864 First released in Issue 3.

41865 Entry included for alignment with the POSIX.1-1988 standard.

#### 41866 Issue 4

41867 The *raise()* and *signal()* functions are added to the list of functions that are either reentrant or not  
41868 interruptible by signals; *fpathconf()* is also added to this list and marked as an extension; *ustat()*  
41869 is removed from the list, as this function is withdrawn from the interface definition. It is no  
41870 longer specified whether *abort()*, *exit()*, and *longjmp()* also fall into this category of functions.

41871 The APPLICATION USAGE section is added. Most of this text is moved from the  
41872 DESCRIPTION in Issue 3.

41873 The FUTURE DIRECTIONS section is added.

41874 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 41875 • The type of argument *act* is changed from **struct sigaction\*** to **const struct sigaction\***.
- 41876 • A statement is added to the DESCRIPTION indicating that the consequence of attempting to  
41877 set SIG\_DFL for a signal that cannot be caught or ignored is unspecified. The [EINVAL] error,  
41878 describing one possible reaction to this condition, is added to the ERRORS section.

#### 41879 Issue 4, Version 2

41880 The following changes are incorporated for X/OPEN UNIX conformance:

- 41881 • The DESCRIPTION describes *sa\_sigaction*, the member of the **sigaction** structure that is the  
41882 signal-catching function.
- 41883 • The DESCRIPTION describes the SA\_ONSTACK, SA\_RESETHAND, SA\_RESTART,  
41884 SA\_SIGINFO, SA\_NOCLDWAIT, and SA\_NODEFER settings of *sa\_flags*. The text describes  
41885 the implications of the use of SA\_SIGINFO for the number of arguments passed to the  
41886 signal-catching function. The text also describes the effects of the SA\_NODEFER and  
41887 SA\_RESETHAND flags on the delivery of a signal and on the permanence of an installed  
41888 action.
- 41889 • The DESCRIPTION specifies the effect if the action for the SIGCHLD signal is set to  
41890 SIG\_IGN.
- 41891 • In the DESCRIPTION, additional text describes the effect if the action is a pointer to a  
41892 function. A new bullet covers the case where SA\_SIGINFO is set. SIGBUS is given as an  
41893 additional signal for which the behavior of a process is undefined following a normal return  
41894 from the signal-catching function.

- 41895           • The APPLICATION USAGE section is updated to describe use of an alternate signal stack;  
41896           resumption of the process receiving the signal; coding for compatibility with POSIX.4-1993;  
41897           and implementation of signal-handling functions in BSD.
- 41898 **Issue 5**
- 41899           The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and POSIX  
41900           Threads Extension.
- 41901           In the DESCRIPTION, the second argument to *func* when SA\_SIGINFO is set is no longer  
41902           permitted to be NULL, and the description of permitted **siginfo\_t** contents is expanded by  
41903           reference to <**signal.h**>.
- 41904           Because the X/OPEN UNIX Extension functionality is now folded into the BASE, the  
41905           [ENOTSUP] error is deleted.
- 41906 **Issue 6**
- 41907           The Open Group corrigenda item U028/7 has been applied. In the paragraph entitled “Signal  
41908           Effects on Other Functions”, a reference to *sigpending()* is added.
- 41909           In the DESCRIPTION, the text “Signal Generation and Delivery” is moved to a separate section  
41910           of this volume of IEEE Std. 1003.1-200x.
- 41911           Text describing functionality from the Realtime Signals option is marked.
- 41912           The following changes are made for alignment with the ISO POSIX-1: 1996 standard:
- 41913           • The [ENOTSUP] error condition is added.
- 41914           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 41915           The **restrict** keyword is added to the *sigaction()* prototype for alignment with the  
41916           ISO/IEC 9899:1999 standard.
- 41917           References to the *wait3()* function are removed.

41918 **NAME**

41919 sigaddset — add a signal to a signal set

41920 **SYNOPSIS**

41921 #include <signal.h>

41922 int sigaddset(sigset\_t \*set, int signo);

41923 **DESCRIPTION**

41924 The *sigaddset()* function adds the individual signal specified by the *signo* to the signal set pointed  
41925 to by *set*.

41926 Applications shall call either *sigemptyset()* or *sigfillset()* at least once for each object of type  
41927 **sigset\_t** prior to any other use of that object. If such an object is not initialized in this way, but is  
41928 nonetheless supplied as an argument to any of *sigaction()*, *sigaddset()*, *sigdelset()*, *sigismember()*,  
41929 *sigpending()*, or *sigprocmask()*, the results are undefined.

41930 **RETURN VALUE**

41931 Upon successful completion, *sigaddset()* shall return 0; otherwise, it shall return -1 and set *errno*  
41932 to indicate the error.

41933 **ERRORS**

41934 The *sigaddset()* function may fail if:

41935 [EINVAL] The value of the *signo* argument is an invalid or unsupported signal number. |

41936 **EXAMPLES**

41937 None.

41938 **APPLICATION USAGE**

41939 None.

41940 **RATIONALE**

41941 None.

41942 **FUTURE DIRECTIONS**

41943 None.

41944 **SEE ALSO**

41945 Section 2.4 (on page 528), *sigaction()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*,  
41946 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
41947 <signal.h>

41948 **CHANGE HISTORY**

41949 First released in Issue 3.

41950 Entry included for alignment with the POSIX.1-1988 standard.

41951 **Issue 4**

41952 The word “will” is replaced by the word “may” in the ERRORS section.

41953 **Issue 5**

41954 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in  
41955 previous issues.

41956 **Issue 6**

41957 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

## 41958 NAME

41959 sigaltstack — set and get signal alternate stack context

## 41960 SYNOPSIS

41961 XSI `#include <signal.h>`41962 `int sigaltstack(const stack_t *restrict ss, stack_t *restrict oss);`

41963

## 41964 DESCRIPTION

41965 The *sigaltstack()* function allows a process to define and examine the state of an alternate stack  
 41966 for signal handlers. Signals that have been explicitly declared to execute on the alternate stack  
 41967 shall be delivered on the alternate stack.

41968 If *ss* is not a null pointer, it points to a **stack\_t** structure that specifies the alternate signal stack  
 41969 that shall take effect upon return from *sigaltstack()*. The *ss\_flags* member specifies the new stack  
 41970 state. If it is set to *SS\_DISABLE*, the stack is disabled and *ss\_sp* and *ss\_size* are ignored.  
 41971 Otherwise, the stack shall be enabled, and the *ss\_sp* and *ss\_size* members specify the new address  
 41972 and size of the stack.

41973 The range of addresses starting at *ss\_sp* up to but not including *ss\_sp+ss\_size*, is available to the  
 41974 implementation for use as the stack. This function makes no assumptions regarding which end  
 41975 is the stack base and in which direction the stack grows as items are pushed.

41976 If *oss* is not a null pointer, on successful completion it shall point to a **stack\_t** structure that  
 41977 specifies the alternate signal stack that was in effect prior to the call to *sigaltstack()*. The *ss\_sp*  
 41978 and *ss\_size* members specify the address and size of that stack. The *ss\_flags* member specifies the  
 41979 stack's state, and may contain one of the following values:

41980 **SS\_ONSTACK** The process is currently executing on the alternate signal stack. Attempts to  
 41981 modify the alternate signal stack while the process is executing on it fail. This  
 41982 flag shall not be modified by processes.

41983 **SS\_DISABLE** The alternate signal stack is currently disabled.

41984 The value *SIGSTKSZ* is a system default specifying the number of bytes that would be used to  
 41985 cover the usual case when manually allocating an alternate stack area. The value *MINSIGSTKSZ*  
 41986 is defined to be the minimum stack size for a signal handler. In computing an alternate stack  
 41987 size, a program should add that amount to its stack requirements to allow for the system  
 41988 implementation overhead. The constants *SS\_ONSTACK*, *SS\_DISABLE*, *SIGSTKSZ*, and  
 41989 *MINSIGSTKSZ* are defined in **<signal.h>**.

41990 After a successful call to one of the *exec* functions, there are no alternate signal stacks in the new  
 41991 process image.

41992 In some implementations, a signal (whether or not indicated to execute on the alternate stack)  
 41993 shall always execute on the alternate stack if it is delivered while another signal is being caught  
 41994 using the alternate stack.

41995 Use of this function by library threads that are not bound to kernel-scheduled entities results in  
 41996 undefined behavior.

## 41997 RETURN VALUE

41998 Upon successful completion, *sigaltstack()* shall return 0; otherwise, it shall return *-1* and set *errno*  
 41999 to indicate the error.

**42000 ERRORS**

42001 The *sigaltstack()* function shall fail if:

42002 [EINVAL] The *ss* argument is not a null pointer, and the *ss\_flags* member pointed to by *ss* |  
 42003 contains flags other than *SS\_DISABLE*.

42004 [ENOMEM] The size of the alternate stack area is less than *MINSIGSTKSZ*. |

42005 [EPERM] An attempt was made to modify an active stack. |

**42006 EXAMPLES****42007 Allocating Memory for an Alternate Stack**

42008 The following example illustrates a method for allocating memory for an alternate stack.

```
42009 #include <signal.h>
42010 ...
42011 if ((sigstk.ss_sp = malloc(SIGSTKSZ)) == NULL)
42012 /* Error return. */
42013 sigstk.ss_size = SIGSTKSZ;
42014 sigstk.ss_flags = 0;
42015 if (sigaltstack(&sigstk, (stack_t *)0) < 0)
42016 perror("sigaltstack");
```

**42017 APPLICATION USAGE**

42018 On some implementations, stack space is automatically extended as needed. On those  
 42019 implementations, automatic extension is typically not available for an alternate stack. If the stack  
 42020 overflows, the behavior is undefined.

**42021 RATIONALE**

42022 None.

**42023 FUTURE DIRECTIONS**

42024 None.

**42025 SEE ALSO**

42026 Section 2.4 (on page 528), *sigaction()*, *sigsetjmp()*, the Base Definitions volume of |  
 42027 IEEE Std. 1003.1-200x, <**signal.h**>

**42028 CHANGE HISTORY**

42029 First released in Issue 4, Version 2.

**42030 Issue 5**

42031 Moved from X/OPEN UNIX extension to BASE.

42032 The last sentence of the DESCRIPTION was included as an APPLICATION USAGE note in |  
 42033 previous issues.

**42034 Issue 6**

42035 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

42036 The **restrict** keyword is added to the *sigaltstack()* prototype for alignment with the |  
 42037 ISO/IEC 9899:1999 standard.

42038 **NAME**

42039 sigdelset — delete a signal from a signal set

42040 **SYNOPSIS**

42041 #include <signal.h>

42042 int sigdelset(sigset\_t \*set, int signo);

42043 **DESCRIPTION**

42044 The *sigdelset()* function deletes the individual signal specified by *signo* from the signal set pointed to by *set*.

42046 Applications should call either *sigemptyset()* or *sigfillset()* at least once for each object of type **sigset\_t** prior to any other use of that object. If such an object is not initialized in this way, but is nonetheless supplied as an argument to any of *sigaction()*, *sigaddset()*, *sigdelset()*, *sigismember()*, *sigpending()*, or *sigprocmask()*, the results are undefined.

42050 **RETURN VALUE**

42051 Upon successful completion, *sigdelset()* shall return 0; otherwise, it shall return -1 and set *errno* to indicate the error.

42053 **ERRORS**

42054 The *sigdelset()* function may fail if:

42055 [EINVAL] The *signo* argument is not a valid signal number, or is an unsupported signal number.

42057 **EXAMPLES**

42058 None.

42059 **APPLICATION USAGE**

42060 None.

42061 **RATIONALE**

42062 None.

42063 **FUTURE DIRECTIONS**

42064 None.

42065 **SEE ALSO**

42066 Section 2.4 (on page 528), *sigaction()*, *sigaddset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*, *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <signal.h>

42069 **CHANGE HISTORY**

42070 First released in Issue 3.

42071 Entry included for alignment with the POSIX.1-1988 standard.

42072 **Issue 4**

42073 The word “will” is replaced by the word “may” in the ERRORS section.

42074 **Issue 5**

42075 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in previous issues.

42077 **NAME**

42078 sigemptyset — initialize and empty a signal set

42079 **SYNOPSIS**

42080 #include <signal.h>

42081 int sigemptyset(sigset\_t \*set);

42082 **DESCRIPTION**

42083 The *sigemptyset()* function initializes the signal set pointed to by *set*, such that all signals defined  
42084 in this volume of IEEE Std. 1003.1-200x are excluded.

42085 **RETURN VALUE**

42086 Upon successful completion, *sigemptyset()* shall return 0; otherwise, it shall return -1 and set  
42087 *errno* to indicate the error.

42088 **ERRORS**

42089 No errors are defined.

42090 **EXAMPLES**

42091 None.

42092 **APPLICATION USAGE**

42093 None.

42094 **RATIONALE**

42095 The implementation of the *sigemptyset()* (or *sigfillset()*) function could quite trivially clear (or  
42096 set) all the bits in the signal set. Alternatively, it would be reasonable to initialize part of the  
42097 structure, such as a version field, to permit binary-compatibility between releases where the size  
42098 of the set varies. For such reasons, either *sigemptyset()* or *sigfillset()* must be called prior to any  
42099 other use of the signal set, even if such use is read-only (for example, as an argument to  
42100 *sigpending()*). This function is not intended for dynamic allocation.

42101 The *sigfillset()* and *sigemptyset()* functions require that the resulting signal set include (or  
42102 exclude) all the signals defined in this volume of IEEE Std. 1003.1-200x. Although it is outside  
42103 the scope of this volume of IEEE Std. 1003.1-200x to place this requirement on signals that are  
42104 implemented as extensions, it is recommended that implementation-defined signals also be  
42105 affected by these functions. However, there may be a good reason for a particular signal not to  
42106 be affected. For example, blocking or ignoring an implementation-defined signal may have  
42107 undesirable side effects, whereas the default action for that signal is harmless. In such a case, it  
42108 would be preferable for such a signal to be excluded from the signal set returned by *sigfillset()*.

42109 In early proposals there was no distinction between invalid and unsupported signals (the names  
42110 of optional signals that were not supported by an implementation were not defined by that  
42111 implementation). The [EINVAL] error was thus specified as a required error for invalid signals.  
42112 With that distinction, it is not necessary to require implementations of these functions to  
42113 determine whether an optional signal is actually supported, as that could have a significant  
42114 performance impact for little value. The error could have been required for invalid signals and  
42115 optional for unsupported signals, but this seemed unnecessarily complex. Thus, the error is  
42116 optional in both cases.

42117 **FUTURE DIRECTIONS**

42118 None.

42119 **SEE ALSO**

42120 Section 2.4 (on page 528), *sigaction()*, *sigaddset()*, *sigdelset()*, *sigfillset()*, *sigismember()*,  
42121 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
42122 <signal.h>

42123 **CHANGE HISTORY**

42124           First released in Issue 3.

42125           Entry included for alignment with the POSIX.1-1988 standard.

42126 **NAME**

42127 sigfillset — initialize and fill a signal set

42128 **SYNOPSIS**

42129 #include <signal.h>

42130 int sigfillset(sigset\_t \*set);

42131 **DESCRIPTION**

42132 The *sigfillset()* function initializes the signal set pointed to by *set*, such that all signals defined in  
42133 this volume of IEEE Std. 1003.1-200x are included.

42134 **RETURN VALUE**

42135 Upon successful completion, *sigfillset()* shall return 0; otherwise, it shall return -1 and set *errno*  
42136 to indicate the error.

42137 **ERRORS**

42138 No errors are defined.

42139 **EXAMPLES**

42140 None.

42141 **APPLICATION USAGE**

42142 None.

42143 **RATIONALE**

42144 Refer to *sigemptyset()* (on page 1863).

42145 **FUTURE DIRECTIONS**

42146 None.

42147 **SEE ALSO**

42148 Section 2.4 (on page 528), *sigaction()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigismember()*,  
42149 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
42150 <signal.h>

42151 **CHANGE HISTORY**

42152 First released in Issue 3.

42153 Entry included for alignment with the POSIX.1-1988 standard.

42154 **NAME**

42155           sighold, sigignore — add a signal to the signal mask or set a signal disposition to be ignored

42156 **SYNOPSIS**

```
42157 XSI #include <signal.h>
```

```
42158 int sighold(int sig);
```

```
42159 int sigignore(int sig);
```

42160

42161 **DESCRIPTION**

42162           Refer to *signal()*.

42163 **NAME**

42164 siginterrupt — allow signals to interrupt functions

42165 **SYNOPSIS**

42166 XSI #include &lt;signal.h&gt;

42167 int siginterrupt(int *sig*, int *flag*);

42168

42169 **DESCRIPTION**42170 The *siginterrupt()* function is used to change the restart behavior when a function is interrupted  
42171 by the specified signal. The function *siginterrupt(sig, flag)* has an effect as if implemented as:

```

42172 siginterrupt(int sig, int flag) {
42173 int ret;
42174 struct sigaction act;
42175 (void) sigaction(sig, NULL, &act);
42176 if (flag)
42177 act.sa_flags &= ~SA_RESTART;
42178 else
42179 act.sa_flags |= SA_RESTART;
42180 ret = sigaction(sig, &act, NULL);
42181 return ret;
42182 }
```

42183 **RETURN VALUE**42184 Upon successful completion, *siginterrupt()* shall return 0; otherwise, -1 shall be returned and  
42185 *errno* set to indicate the error.42186 **ERRORS**42187 The *siginterrupt()* function shall fail if:42188 [EINVAL] The *sig* argument is not a valid signal number.42189 **EXAMPLES**

42190 None.

42191 **APPLICATION USAGE**42192 The *siginterrupt()* function supports programs written to historical system interfaces. A portable  
42193 application, when being written or rewritten, should use *sigaction()* with the SA\_RESTART flag  
42194 instead of *siginterrupt()*.42195 **RATIONALE**

42196 None.

42197 **FUTURE DIRECTIONS**

42198 None.

42199 **SEE ALSO**42200 Section 2.4 (on page 528), *sigaction()*, the Base Definitions volume of IEEE Std. 1003.1-200x,

42201 &lt;signal.h&gt;

42202 **CHANGE HISTORY**

42203 First released in Issue 4, Version 2.

42204 **Issue 5**

42205 Moved from X/OPEN UNIX extension to BASE.

42206 **NAME**

42207 sigismember — test for a signal in a signal set

42208 **SYNOPSIS**

42209 #include &lt;signal.h&gt;

42210 int sigismember(const sigset\_t \*set, int signo);

42211 **DESCRIPTION**42212 The *sigismember()* function tests whether the signal specified by *signo* is a member of the set  
42213 pointed to by *set*.42214 Applications should call either *sigemptyset()* or *sigfillset()* at least once for each object of type  
42215 **sigset\_t** prior to any other use of that object. If such an object is not initialized in this way, but is  
42216 nonetheless supplied as an argument to any of *sigaction()*, *sigaddset()*, *sigdelset()*, *sigismember()*,  
42217 *sigpending()*, or *sigprocmask()*, the results are undefined.42218 **RETURN VALUE**42219 Upon successful completion, *sigismember()* shall return 1 if the specified signal is a member of  
42220 the specified set, or 0 if it is not. Otherwise, it shall return -1 and set *errno* to indicate the error.42221 **ERRORS**42222 The *sigismember()* function may fail if:42223 [EINVAL] The *signo* argument is not a valid signal number, or is an unsupported signal  
42224 number.42225 **EXAMPLES**

42226 None.

42227 **APPLICATION USAGE**

42228 None.

42229 **RATIONALE**

42230 None.

42231 **FUTURE DIRECTIONS**

42232 None.

42233 **SEE ALSO**42234 Section 2.4 (on page 528), *sigaction()*, *sigaddset()*, *sigdelset()*, *sigfillset()*, *sigemptyset()*,  
42235 *sigpending()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
42236 <signal.h>42237 **CHANGE HISTORY**

42238 First released in Issue 3.

42239 Entry included for alignment with the POSIX.1-1988 standard.

42240 **Issue 4**

42241 The following changes are incorporated for alignment with the ISO C standard:

- 42242
- The type of the argument *set* is changed from **sigset\_t\*** to type **const sigset\_t\***.
  - The word “will” is replaced by the word “may” in the ERRORS section.
- 42243

42244 **Issue 5**42245 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in  
42246 previous issues.

42247 **NAME**

42248 siglongjmp — non-local goto with signal handling

42249 **SYNOPSIS**

42250 #include &lt;setjmp.h&gt;

42251 void siglongjmp(sigjmp\_buf env, int val);

42252 **DESCRIPTION**

42253 The *siglongjmp()* function restores the environment saved by the most recent invocation of  
42254 *sigsetjmp()* in the same thread, with the corresponding **sigjmp\_buf** argument. If there is no such  
42255 invocation, or if the function containing the invocation of *sigsetjmp()* has terminated execution in  
42256 the interim, the behavior is undefined.

42257 All accessible objects have values as of the time *siglongjmp()* was called, except that the values of  
42258 objects of automatic storage duration which are local to the function containing the invocation of  
42259 the corresponding *sigsetjmp()* which do not have volatile-qualified type and which are changed  
42260 between the *sigsetjmp()* invocation and *siglongjmp()* call are indeterminate.

42261 As it bypasses the usual function call and return mechanisms, *siglongjmp()* shall execute  
42262 correctly in contexts of interrupts, signals, and any of their associated functions. However, if  
42263 *siglongjmp()* is invoked from a nested signal handler (that is, from a function invoked as a result  
42264 of a signal raised during the handling of another signal), the behavior is undefined.

42265 The *siglongjmp()* function shall restore the saved signal mask if and only if the *env* argument was  
42266 initialized by a call to *sigsetjmp()* with a non-zero *savemask* argument.

42267 The effect of a call to *siglongjmp()* where initialization of the **jmp\_buf** structure was not  
42268 performed in the calling thread is undefined.

42269 **RETURN VALUE**

42270 After *siglongjmp()* is completed, program execution shall continue as if the corresponding  
42271 invocation of *sigsetjmp()* had just returned the value specified by *val*. The *siglongjmp()* function  
42272 shall not cause *sigsetjmp()* to return 0; if *val* is 0, *sigsetjmp()* shall return the value 1.

42273 **ERRORS**

42274 No errors are defined.

42275 **EXAMPLES**

42276 None.

42277 **APPLICATION USAGE**

42278 The distinction between *setjmp()* or *longjmp()* and *sigsetjmp()* or *siglongjmp()* is only significant  
42279 for programs which use *sigaction()*, *sigprocmask()*, or *sigsuspend()*.

42280 **RATIONALE**

42281 None.

42282 **FUTURE DIRECTIONS**

42283 None.

42284 **SEE ALSO**

42285 *longjmp()*, *setjmp()*, *sigprocmask()*, *sigsetjmp()*, *sigsuspend()*, the Base Definitions volume of  
42286 IEEE Std. 1003.1-200x, <**setjmp.h**>

42287 **CHANGE HISTORY**

42288 First released in Issue 3.

42289 Entry included for alignment with the ISO POSIX-1 standard.

42290 **Issue 4**

42291 The APPLICATION USAGE section is amended.

42292 An ERRORS section is added.

42293 **Issue 5**

42294 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

## 42295 NAME

42296 sighold, sigignore, signal, sigpause, sigrelse, sigset — signal management

## 42297 SYNOPSIS

42298 #include &lt;signal.h&gt;

42299 void (\*signal(int *sig*, void (\**func*)(int)))(int);42300 XSI int sighold(int *sig*);42301 int sigignore(int *sig*);42302 int sigpause(int *sig*);42303 int sigrelse(int *sig*);42304 void (\*sigset(int *sig*, void (\**disp*)(int)))(int);

42305

## 42306 DESCRIPTION

42307 CX For *signal()*: The functionality described on this reference page is aligned with the ISO C  
 42308 standard. Any conflict between the requirements described here and the ISO C standard is  
 42309 unintentional. This volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

42310 CX Use of any of these functions is unspecified in a multi-threaded process.

42311 The *signal()* function chooses one of three ways in which receipt of the signal number *sig* is to be  
 42312 subsequently handled. If the value of *func* is SIG\_DFL, default handling for that signal shall  
 42313 occur. If the value of *func* is SIG\_IGN, the signal shall be ignored. Otherwise, the application  
 42314 shall ensure that *func* points to a function to be called when that signal occurs. An invocation of  
 42315 such a function because of a signal, or (recursively) of any further functions called by that  
 42316 invocation (other than functions in the standard library), is called a “signal handler”.

42317 When a signal occurs, if *func* points to a function, first the equivalent of a:

42318 `signal(sig, SIG_DFL);`

42319 is executed or an implementation-defined blocking of the signal is performed. (If the value of *sig*  
 42320 is SIGILL, whether the reset to SIG\_DFL occurs is implementation-defined.) Next the equivalent  
 42321 of:

42322 `(*func)(sig);`

42323 is executed. The *func* function may terminate by executing a **return** statement or by calling  
 42324 *abort()*, *exit()*, or *longjmp()*. If *func* executes a **return** statement and the value of *sig* was SIGFPE  
 42325 or any other implementation-defined value corresponding to a computational exception, the  
 42326 behavior is undefined. Otherwise, the program shall resume execution at the point it was  
 42327 interrupted. If the signal occurs as the result of calling the *abort()* or *raise()* function, the signal  
 42328 handler shall not call the *raise()* function.

42329 If the signal occurs other than as the result of calling *abort()*, *kill()*, or *raise()*, the behavior is  
 42330 undefined if the signal handler refers to any object with static storage duration other than by  
 42331 assigning a value to an object declared as volatile **sig\_atomic\_t**, or if the signal handler calls any  
 42332 function in the standard library other than one of the functions listed in Section 2.4 (on page 528)  
 42333 or refers to any object with static storage duration other than by assigning a value to a static  
 42334 storage duration variable of type volatile **sig\_atomic\_t**. Furthermore, if such a call fails, the  
 42335 value of *errno* is indeterminate.

42336 At program start-up, the equivalent of:

42337 `signal(sig, SIG_IGN);`

42338 is executed for some signals, and the equivalent of:

- 42339 `signal(sig, SIG_DFL);`
- 42340 is executed for all other signals (see *exec*).
- 42341 XSI The *sighold()*, *sigignore()*, *sigpause()*, *sigrelse()*, and *sigset()* functions provide simplified signal  
42342 management.
- 42343 The *sigset()* function is used to modify signal dispositions. The *sig* argument specifies the signal,  
42344 which may be any signal except SIGKILL and SIGSTOP. The *disp* argument specifies the signal's  
42345 disposition, which may be SIG\_DFL, SIG\_IGN, or the address of a signal handler. If *sigset()* is  
42346 used, and *disp* is the address of a signal handler, the system shall add *sig* to the calling process'  
42347 signal mask before executing the signal handler; when the signal handler returns, the system  
42348 shall restore the calling process' signal mask to its state prior the delivery of the signal. In  
42349 addition, if *sigset()* is used, and *disp* is equal to SIG\_HOLD, *sig* shall be added to the calling  
42350 process' signal mask and *sig*'s disposition shall remain unchanged. If *sigset()* is used, and *disp* is  
42351 not equal to SIG\_HOLD, *sig* shall be removed from the calling process' signal mask.
- 42352 The *sighold()* function adds *sig* to the calling process' signal mask.
- 42353 The *sigrelse()* function removes *sig* from the calling process' signal mask.
- 42354 The *sigignore()* function sets the disposition of *sig* to SIG\_IGN.
- 42355 The *sigpause()* function removes *sig* from the calling process' signal mask and suspends the  
42356 calling process until a signal is received. The *sigpause()* function restores the process' signal  
42357 mask to its original state before returning.
- 42358 If the action for the SIGCHLD signal is set to SIG\_IGN, child processes of the calling processes  
42359 shall not be transformed into zombie processes when they terminate. If the calling process  
42360 subsequently waits for its children, and the process has no unwaited-for children that were  
42361 transformed into zombie processes, it shall block until all of its children terminate, and *wait()*,  
42362 *waitid()*, and *waitpid()* shall fail and set *errno* to [ECHILD].
- 42363 **RETURN VALUE**
- 42364 If the request can be honored, *signal()* shall return the value of *func* for the most recent call to  
42365 *signal()* for the specified signal *sig*. Otherwise, SIG\_ERR shall be returned and a positive value  
42366 shall be stored in *errno*.
- 42367 XSI Upon successful completion, *sigset()* shall return SIG\_HOLD if the signal had been blocked and  
42368 the signal's previous disposition if it had not been blocked. Otherwise, SIG\_ERR shall be  
42369 returned and *errno* set to indicate the error.
- 42370 The *sigpause()* function shall suspend execution of the thread until a signal is received,  
42371 whereupon it shall return -1 and set *errno* to [EINTR].
- 42372 For all other functions, upon successful completion, 0 shall be returned. Otherwise, -1 shall be  
42373 returned and *errno* set to indicate the error.
- 42374 **ERRORS**
- 42375 The *signal()* function shall fail if:
- 42376 CX [EINVAL] The *sig* argument is not a valid signal number or an attempt is made to catch a  
42377 signal that cannot be caught or ignore a signal that cannot be ignored.
- 42378 The *signal()* function may fail if:
- 42379 CX [EINVAL] An attempt was made to set the action to SIG\_DFL for a signal that cannot be  
42380 caught or ignored (or both).
- 42381 The *sighold()*, *sigignore()*, *sigpause()*, *sigrelse()*, and *sigset()* functions shall fail if:

- 42382 XSI [EINVAL] The *sig* argument is an illegal signal number.
- 42383 The *sigset()* and *sigignore()* functions shall fail if:
- 42384 XSI [EINVAL] An attempt is made to catch a signal that cannot be caught, or to ignore a  
42385 signal that cannot be ignored.
- 42386 **EXAMPLES**
- 42387 None.
- 42388 **APPLICATION USAGE**
- 42389 The *sigaction()* function provides a more comprehensive and reliable mechanism for controlling  
42390 signals; new applications should use *sigaction()* rather than *signal()*.
- 42391 The *sighold()* function, in conjunction with *sigrelse()* or *sigpause()*, may be used to establish  
42392 critical regions of code that require the delivery of a signal to be temporarily deferred.
- 42393 The *sigsuspend()* function should be used in preference to *sigpause()* for broader portability.
- 42394 **RATIONALE**
- 42395 None.
- 42396 **FUTURE DIRECTIONS**
- 42397 None.
- 42398 **SEE ALSO**
- 42399 Section 2.4 (on page 528), *exec*, *pause()*, *sigaction()*, *sigsuspend()*, *waitid()*, the Base Definitions  
42400 volume of IEEE Std. 1003.1-200x, <signal.h>
- 42401 **CHANGE HISTORY**
- 42402 First released in Issue 1. Derived from Issue 1 of the SVID.
- 42403 **Issue 4**
- 42404 The APPLICATION USAGE section is added.
- 42405 The following changes are incorporated for alignment with the ISO C standard:
- 42406 • The function is no longer marked as an extension.
- 42407 • The argument **int** is added to the definition of *func* in the SYNOPSIS section.
- 42408 • In Issue 3, this function cross-referred to *sigaction()*. This issue provides a complete  
42409 description of the function as defined in ISO C standard.
- 42410 **Issue 4, Version 2**
- 42411 The following changes are incorporated for X/OPEN UNIX conformance:
- 42412 • The *sighold()*, *sigignore()*, *sigpause()*, *sigrelse()*, and *sigset()* functions are added to the  
42413 SYNOPSIS.
- 42414 • The DESCRIPTION is updated to describe semantics of the above functions.
- 42415 • Additional text is added to the RETURN VALUE section to describe possible returns from  
42416 the *sigset()* function specifically, and all of the above functions in general.
- 42417 • The ERRORS section is restructured to describe possible error returns from each of the above  
42418 functions individually.
- 42419 • The APPLICATION USAGE section is updated to describe certain programming  
42420 considerations associated with the X/OPEN UNIX functions.

42421 **Issue 5**

42422 Moved from X/OPEN UNIX extension to BASE.

42423 The DESCRIPTION is updated to indicate that the *sigpause()* function restores the process' signal mask to its original state before returning.

42425 The RETURN VALUE section is updated to indicate that the *sigpause()* function suspends execution of the process until a signal is received, whereupon it returns -1 and sets *errno* to [EINTR].

42428 **Issue 6**

42429 Extensions beyond the ISO C standard are now marked.

42430 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

42431 The DESCRIPTION is updated for alignment with the ISO/IEC 9899:1999 standard.

42432 References to the *wait3()* function are removed.

42433 **NAME**

42434 signbit — test sign

42435 **SYNOPSIS**

42436 #include &lt;math.h&gt;

42437 int signbit(real-floating x);

42438 **DESCRIPTION**

42439 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
42440 conflict between the requirements described here and the ISO C standard is unintentional. This  
42441 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

42442 The *signbit()* macro shall determine whether the sign of its argument value is negative.42443 **RETURN VALUE**

42444 The *signbit()* macro shall return a non-zero value if and only if the sign of its argument value is  
42445 negative.

42446 **ERRORS**

42447 No errors are defined.

42448 **EXAMPLES**

42449 None.

42450 **APPLICATION USAGE**

42451 None.

42452 **RATIONALE**

42453 None.

42454 **FUTURE DIRECTIONS**

42455 None.

42456 **SEE ALSO**

42457 *fpclassify()*, *isfinite()*, *isinf()*, *isnan()*, *isnormal()*, the Base Definitions volume of  
42458 IEEE Std. 1003.1-200x, <math.h>

42459 **CHANGE HISTORY**

42460 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

42461 **NAME**

42462 sigpause — remove a signal from the signal mask and suspend the thread

42463 **SYNOPSIS**

```
42464 xSI #include <signal.h>
```

```
42465 int sigpause(int sig);
```

42466

42467 **DESCRIPTION**

42468 Refer to *signal()*.

42469 **NAME**

42470 sigpending — examine pending signals

42471 **SYNOPSIS**

42472 #include <signal.h>

42473 int sigpending(sigset\_t \*set);

42474 **DESCRIPTION**

42475 The *sigpending()* function stores, in the location referenced by the *set* argument, the set of signals  
42476 that are blocked from delivery to the calling thread and that are pending on the process or the  
42477 calling thread.

42478 **RETURN VALUE**

42479 Upon successful completion, *sigpending()* shall return 0; otherwise, -1 shall be returned and  
42480 *errno* set to indicate the error.

42481 **ERRORS**

42482 No errors are defined.

42483 **EXAMPLES**

42484 None.

42485 **APPLICATION USAGE**

42486 None.

42487 **RATIONALE**

42488 None.

42489 **FUTURE DIRECTIONS**

42490 None.

42491 **SEE ALSO**

42492 *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*, *sigprocmask()*, the Base Definitions  
42493 volume of IEEE Std. 1003.1-200x, <signal.h>

42494 **CHANGE HISTORY**

42495 First released in Issue 3.

42496 **Issue 5**

42497 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

42498 **NAME**

42499 sigprocmask — examine and change blocked signals

42500 **SYNOPSIS**

42501 #include &lt;signal.h&gt;

42502 int sigprocmask(int *how*, const sigset\_t \*restrict *set*,  
42503 sigset\_t \*restrict *oset*);42504 **DESCRIPTION**42505 Refer to *pthread\_sigmask()*.

## 42506 NAME

42507 sigqueue — queue a signal to a process (**REALTIME**)

## 42508 SYNOPSIS

42509 RTS `#include <signal.h>`42510 `int sigqueue(pid_t pid, int signo, const union sigval value);`

42511

## 42512 DESCRIPTION

42513 The *sigqueue()* function shall cause the signal specified by *signo* to be sent with the value  
 42514 specified by *value* to the process specified by *pid*. If *signo* is zero (the null signal), error checking  
 42515 is performed but no signal is actually sent. The null signal can be used to check the validity of  
 42516 *pid*.

42517 The conditions required for a process to have permission to queue a signal to another process  
 42518 are the same as for the *kill()* function.

42519 The *sigqueue()* function returns immediately. If SA\_SIGINFO is set for *signo* and if the resources  
 42520 were available to queue the signal, the signal is queued and sent to the receiving process. If  
 42521 SA\_SIGINFO is not set for *signo*, then *signo* is sent at least once to the receiving process; it is  
 42522 unspecified whether *value* shall be sent to the receiving process as a result of this call.

42523 If the value of *pid* causes *signo* to be generated for the sending process, and if *signo* is not blocked  
 42524 for the calling thread and if no other thread has *signo* unblocked or is waiting in a *sigwait()*  
 42525 function for *signo*, either *signo* or at least the pending, unblocked signal shall be delivered to the  
 42526 calling thread before the *sigqueue()* function returns. Should any multiple pending signals in the  
 42527 range SIGRTMIN to SIGRTMAX be selected for delivery, it shall be the lowest numbered one.  
 42528 The selection order between realtime and non-realtime signals, or between multiple pending  
 42529 non-realtime signals, is unspecified.

## 42530 RETURN VALUE

42531 Upon successful completion, the specified signal shall have been queued, and the *sigqueue()*  
 42532 function shall return a value of zero. Otherwise, the function shall return a value of  $-1$  and set  
 42533 *errno* to indicate the error.

## 42534 ERRORS

42535 The *sigqueue()* function shall fail if:

|       |          |                                                                              |  |
|-------|----------|------------------------------------------------------------------------------|--|
| 42536 | [EAGAIN] | No resources available to queue the signal. The process has already queued   |  |
| 42537 |          | SIGQUEUE_MAX signals that are still pending at the receiver(s), or a system- |  |
| 42538 |          | wide resource limit has been exceeded.                                       |  |

|       |          |                                                                                    |  |
|-------|----------|------------------------------------------------------------------------------------|--|
| 42539 | [EINVAL] | The value of the <i>signo</i> argument is an invalid or unsupported signal number. |  |
|-------|----------|------------------------------------------------------------------------------------|--|

|       |         |                                                                               |  |
|-------|---------|-------------------------------------------------------------------------------|--|
| 42540 | [EPERM] | The process does not have the appropriate privilege to send the signal to the |  |
| 42541 |         | receiving process.                                                            |  |

|       |         |                                        |  |
|-------|---------|----------------------------------------|--|
| 42542 | [ESRCH] | The process <i>pid</i> does not exist. |  |
|-------|---------|----------------------------------------|--|

42543 **EXAMPLES**

42544 None.

42545 **APPLICATION USAGE**

42546 None.

42547 **RATIONALE**

42548 The *sigqueue()* function allows an application to queue a realtime signal to itself or to another  
42549 process, specifying the application-defined value. This is common practice in realtime  
42550 applications on existing realtime systems. It was felt that specifying another function in the  
42551 *sig...* name space already carved out for signals was preferable to extending the function to  
42552 *kill()*.

42553 Such a function became necessary when the put/get event function of the message queues was  
42554 removed. It should be noted that the *sigqueue()* function implies reduced performance in a  
42555 security-conscious implementation as the access permissions between the sender and receiver  
42556 have to be checked on each send when the *pid* is resolved into a target process. Such access  
42557 checks were necessary only at message queue open in the previous function.

42558 The standard developers required that *sigqueue()* have the same semantics with respect to the  
42559 null signal as *kill()*, and that the same permission checking be used. But because of the difficulty  
42560 of implementing the “broadcast” semantic of *kill()* (for example, to process groups) and the  
42561 interaction with resource allocation, this semantic was not adopted. The *sigqueue()* function  
42562 queues a signal to a single process specified by the *pid* argument.

42563 The *sigqueue()* function can fail if the system has insufficient resources to queue the signal. An  
42564 explicit limit on the number of queued signals that a process could send was introduced. While  
42565 the limit is “per-sender”, this volume of IEEE Std. 1003.1-200x does not specify that the  
42566 resources be part of the state of the sender. This would require either that the sender be  
42567 maintained after exit until all signals that it had sent to other processes were handled or that all  
42568 such signals that had not yet been acted upon be removed from the queue(s) of the receivers.  
42569 This volume of IEEE Std. 1003.1-200x does not preclude this behavior, but an implementation  
42570 that allocated queuing resources from a system-wide pool (with per-sender limits) and that  
42571 leaves queued signals pending after the sender exits is also permitted.

42572 **FUTURE DIRECTIONS**

42573 None.

42574 **SEE ALSO**42575 Section 2.8.1 (on page 543), the Base Definitions volume of IEEE Std. 1003.1-200x, <**signal.h**>42576 **CHANGE HISTORY**

42577 First released in Issue 5. Included for alignment with the POSIX Realtime Extension and the  
42578 POSIX Threads Extension.

42579 **Issue 6**42580 The *sigqueue()* function is marked as part of the Realtime Signals Extension option.

42581 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
42582 implementation does not support the Realtime Signals Extension option.

42583 **NAME**

42584 sigrelse, sigset — remove a signal from signal mask or modify signal disposition

42585 **SYNOPSIS**

42586 XSI #include <signal.h>

42587 int sigrelse(int *sig*);

42588 void (\*sigset(int *sig*, void (\**disp*)(int)))(int);

42589

42590 **DESCRIPTION**

42591 Refer to *signal()*.

42592 **NAME**

42593 sigsetjmp — set jump point for a non-local goto

42594 **SYNOPSIS**

42595 #include &lt;setjmp.h&gt;

42596 int sigsetjmp(sigjmp\_buf env, int savemask);

42597 **DESCRIPTION**

42598 A call to *sigsetjmp()* saves the calling environment in its *env* argument for later use by  
 42599 *siglongjmp()*. It is unspecified whether *sigsetjmp()* is a macro or a function. If a macro definition  
 42600 is suppressed in order to access an actual function, or a program defines an external identifier  
 42601 with the name *sigsetjmp*, the behavior is undefined.

42602 If the value of the *savemask* argument is not 0, *sigsetjmp()* shall also save the current signal mask  
 42603 of the calling thread as part of the calling environment.

42604 All accessible objects have values as of the time *siglongjmp()* was called, except that the values of  
 42605 objects of automatic storage duration which are local to the function containing the invocation of  
 42606 the corresponding *sigsetjmp()* which do not have volatile-qualified type and which are changed  
 42607 between the *sigsetjmp()* invocation and *siglongjmp()* call are indeterminate.

42608 The application shall ensure that an invocation of *sigsetjmp()* appears in one of the following  
 42609 contexts only:

- 42610 • The entire controlling expression of a selection or iteration statement
- 42611 • One operand of a relational or equality operator with the other operand an integral constant  
 42612 expression, with the resulting expression being the entire controlling expression of a  
 42613 selection or iteration statement
- 42614 • The operand of a unary ('!') operator with the resulting expression being the entire  
 42615 controlling expression of a selection or iteration
- 42616 • The entire expression of an expression statement (possibly cast to **void**)

42617 **RETURN VALUE**

42618 If the return is from a successful direct invocation, *sigsetjmp()* shall return 0. If the return is from  
 42619 a call to *siglongjmp()*, *sigsetjmp()* shall return a non-zero value.

42620 **ERRORS**

42621 No errors are defined.

42622 **EXAMPLES**

42623 None.

42624 **APPLICATION USAGE**

42625 The distinction between *setjmp()/longjmp()* and *sigsetjmp()/siglongjmp()* is only significant for  
 42626 programs which use *sigaction()*, *sigprocmask()*, or *sigsuspend()*.

42627 **RATIONALE**

42628 The ISO C standard specifies various restrictions on the usage of the *setjmp()* macro in order to  
 42629 permit implementors to recognize the name in the compiler and not implement an actual  
 42630 function. These same restrictions apply to the *sigsetjmp()* macro.

42631 There are processors that cannot easily support these calls, but this was not considered a  
 42632 sufficient reason to exclude them.

42633 4.2 BSD and 4.3 BSD systems provide functions named *\_setjmp()* and *\_longjmp()* that, together  
 42634 with *setjmp()* and *longjmp()*, provide the same functionality as *sigsetjmp()* and *siglongjmp()*. On  
 42635 those systems, *setjmp()* and *longjmp()* save and restore signal masks, while *\_setjmp()* and

42636 *\_longjmp()* do not. On System V, Release 3 and in corresponding issues of the SVID, *setjmp()* and  
42637 *longjmp()* are explicitly defined not to save and restore signal masks. In order to permit existing  
42638 practice in both cases, the relation of *setjmp()* and *longjmp()* to signal masks is not specified, and  
42639 a new set of functions is defined instead.

42640 The *longjmp()* and *siglongjmp()* functions operate as in the previous issue provided the matching  
42641 *setjmp()* or *sigsetjmp()* has been performed in the same thread. Non-local jumps into contexts  
42642 saved by other threads would be at best a questionable practice and were not considered worthy  
42643 of standardization.

#### 42644 FUTURE DIRECTIONS

42645 None.

#### 42646 SEE ALSO

42647 *siglongjmp()*, *signal()*, *sigprocmask()*, *sigsuspend()*, the Base Definitions volume of  
42648 IEEE Std. 1003.1-200x, <setjmp.h>

#### 42649 CHANGE HISTORY

42650 First released in Issue 3.

42651 Entry included for alignment with the POSIX.1-1988 standard.

#### 42652 Issue 4

42653 The DESCRIPTION states that *sigsetjmp()* is a macro or a function. Issue 3 states that it is a  
42654 macro. Warnings are also added about the suppression of a *sigsetjmp()* macro definition.

42655 A statement is added to the DESCRIPTION about the accessibility of objects after a *siglongjmp()*  
42656 call.

42657 Text is added to the DESCRIPTION describing the contexts in which calls to *sigsetjmp()* are  
42658 valid.

#### 42659 Issue 5

42660 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

#### 42661 Issue 6

42662 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

42663 **NAME**

42664 sigsuspend — wait for a signal

42665 **SYNOPSIS**

42666 #include &lt;signal.h&gt;

42667 int sigsuspend(const sigset\_t \*sigmask);

42668 **DESCRIPTION**

42669 The *sigsuspend()* function replaces the current signal mask of the calling thread with the set of  
42670 signals pointed to by *sigmask* and then suspends the thread until delivery of a signal whose  
42671 action is either to execute a signal-catching function or to terminate the process. This shall not  
42672 cause any other signals that may have been pending on the process to become pending on the  
42673 thread.

42674 If the action is to terminate the process then *sigsuspend()* shall never return. If the action is to  
42675 execute a signal-catching function, then *sigsuspend()* shall return after the signal-catching  
42676 function returns, with the signal mask restored to the set that existed prior to the *sigsuspend()*  
42677 call.

42678 It is not possible to block signals that cannot be ignored. This is enforced by the system without  
42679 causing an error to be indicated.

42680 **RETURN VALUE**

42681 Since *sigsuspend()* suspends thread execution indefinitely, there is no successful completion  
42682 return value. If a return occurs, -1 shall be returned and *errno* set to indicate the error.

42683 **ERRORS**42684 The *sigsuspend()* function shall fail if:

42685 [EINTR] A signal is caught by the calling process and control is returned from the  
42686 signal-catching function.

42687 **EXAMPLES**

42688 None.

42689 **APPLICATION USAGE**

42690 Normally, at the beginning of a critical code section, a specified set of signals is blocked using  
42691 the *sigprocmask()* function. When the thread has completed the critical section and needs to wait  
42692 for the previously blocked signal(s), it pauses by calling *sigsuspend()* with the mask that was  
42693 returned by the *sigprocmask()* call.

42694 **RATIONALE**

42695 None.

42696 **FUTURE DIRECTIONS**

42697 None.

42698 **SEE ALSO**

42699 Section 2.4 (on page 528), *pause()*, *sigaction()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, the  
42700 Base Definitions volume of IEEE Std. 1003.1-200x, <signal.h>

42701 **CHANGE HISTORY**

42702 First released in Issue 3.

42703 Entry included for alignment with the POSIX.1-1988 standard.

42704 **Issue 4**

42705 The term “signal handler” is changed to “signal-catching function”.

42706 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 42707 • The type of the argument *sigmask* is changed from **sigset\_t\*** to **const sigset\_t\***.

42708 **Issue 5**

42709 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

42710 **Issue 6**

42711 The text in the RETURN VALUE section has been changed from “suspends process execution”  
42712 to “suspends thread execution”. This reflects IEEE PASC Interpretation 1003.1c #40.

42713 Text in the APPLICATION USAGE section has been replaced.

42714 **NAME**42715 sigtimedwait, sigwaitinfo — wait for queued signals (**REALTIME**)42716 **SYNOPSIS**42717 RTS 

```
#include <signal.h>
```

42718 

```
int sigtimedwait(const sigset_t *restrict set, siginfo_t *restrict info,
```

42719 

```
const struct timespec *restrict timeout);
```

42720 

```
int sigwaitinfo(const sigset_t *restrict set, siginfo_t *restrict info);
```

42721

42722 **DESCRIPTION**

42723 The *sigtimedwait()* function behaves the same as *sigwaitinfo()* except that if none of the signals  
 42724 specified by *set* are pending, *sigtimedwait()* waits for the time interval specified in the **timespec**  
 42725 structure referenced by *timeout*. If the **timespec** structure pointed to by *timeout* is zero-valued  
 42726 and if none of the signals specified by *set* are pending, then *sigtimedwait()* returns immediately  
 42727 with an error. If *timeout* is the NULL pointer, the behavior is unspecified. If the Monotonic Clock  
 42728 option is supported, the **CLOCK\_MONOTONIC** clock shall be used to measure the time interval  
 42729 specified by the *timeout* argument.

42730 The *sigwaitinfo()* function selects the pending signal from the set specified by *set*. Should any of  
 42731 multiple pending signals in the range **SIGRTMIN** to **SIGRTMAX** be selected, it shall be the  
 42732 lowest numbered one. The selection order between realtime and non-realtime signals, or  
 42733 between multiple pending non-realtime signals, is unspecified. If no signal in *set* is pending at  
 42734 the time of the call, the calling thread is suspended until one or more signals in *set* become  
 42735 pending or until it is interrupted by an unblocked, caught signal.

42736 The *sigwaitinfo()* function behaves the same as the *sigwait()* function if the *info* argument is  
 42737 NULL. If the *info* argument is non-NULL, the *sigwaitinfo()* function behaves the same as  
 42738 *sigwait()*, except that the selected signal number shall be stored in the *si\_signo* member, and the  
 42739 cause of the signal shall be stored in the *si\_code* member. If any value is queued to the selected  
 42740 signal, the first such queued value shall be dequeued and, if the *info* argument is non-NULL, the  
 42741 value shall be stored in the *si\_value* member of *info*. The system resource used to queue the  
 42742 signal shall be released and returned to the system for other use. If no value is queued, the  
 42743 content of the *si\_value* member is undefined. If no further signals are queued for the selected  
 42744 signal, the pending indication for that signal shall be reset.

42745 **RETURN VALUE**

42746 Upon successful completion (that is, one of the signals specified by *set* is pending or is  
 42747 generated) *sigwaitinfo()* and *sigtimedwait()* shall return the selected signal number. Otherwise,  
 42748 the function shall return a value of **-1** and set *errno* to indicate the error.

42749 **ERRORS**42750 The *sigtimedwait()* function shall fail if:42751 **Notes to Reviewers**42752 *This section with side shading will not appear in the final copy. - Ed.*

42753 D1, XSH, ERN 345 proposes that the [EAGAIN] error condition ought to be [ETIMEDOUT] as  
 42754 per the same condition on *pthread\_cond\_timedwait()*.

42755 [EAGAIN] No signal specified by *set* was generated within the specified timeout period.

42756 The *sigtimedwait()* and *sigwaitinfo()* functions may fail if:

42757 [EINTR] The wait was interrupted by an unblocked, caught signal. It shall be  
 42758 documented in system documentation whether this error causes these

42759 functions to fail.

42760 The *sigtimedwait()* function may also fail if:

42761 [EINVAL] The *timeout* argument specified a *tv\_nsec* value less than zero or greater than  
42762 or equal to 1 000 million.

42763 An implementation only checks for this error if no signal is pending in *set* and it is necessary to  
42764 wait.

#### 42765 EXAMPLES

42766 None.

#### 42767 APPLICATION USAGE

42768 None.

#### 42769 RATIONALE

42770 Existing programming practice on realtime systems uses the ability to pause waiting for a  
42771 selected set of events and handle the first event that occurs in-line instead of in a signal-handling  
42772 function. This allows applications to be written in an event-directed style similar to a state  
42773 machine. This style of programming is useful for largescale transaction processing in which the  
42774 overall throughput of an application and the ability to clearly track states are more important  
42775 than the ability to minimize the response time of individual event handling.

42776 It is possible to construct a signal-waiting macro function out of the realtime signal function  
42777 mechanism defined in this volume of IEEE Std. 1003.1-200x. However, such a macro has to  
42778 include the definition of a generalized handler for all signals to be waited on. A significant  
42779 portion of the overhead of handler processing can be avoided if the signal-waiting function is  
42780 provided by the kernel. This volume of IEEE Std. 1003.1-200x therefore provides two signal-  
42781 waiting functions—one that waits indefinitely and one with a timeout—as part of the overall  
42782 realtime signal function specification.

42783 The specification of a function with a timeout allows an application to be written that can be  
42784 broken out of a wait after a set period of time if no event has occurred. It was argued that setting  
42785 a timer event before the wait and recognizing the timer event in the wait would also implement  
42786 the same functionality, but at a lower performance level. Because of the performance  
42787 degradation associated with the user-level specification of a timer event and the subsequent  
42788 cancelation of that timer event after the wait completes for a valid event, and the complexity  
42789 associated with handling potential race conditions associated with the user-level method, the  
42790 separate function has been included.

42791 Note that the semantics of the *sigwaitinfo()* function are nearly identical to that of the *sigwait()*  
42792 function defined by this volume of IEEE Std. 1003.1-200x. The only difference is that *sigwaitinfo()*  
42793 returns the queued signal value in the *value* argument. The return of the queued value is  
42794 required so that applications can differentiate between multiple events queued to the same  
42795 signal number.

42796 The two distinct functions are being maintained because some implementations may choose to  
42797 implement the POSIX Threads Extension functions and not implement the queued signals  
42798 extensions. Note, though, that *sigwaitinfo()* does not return the queued value if the *value*  
42799 argument is NULL, so the POSIX Threads Extension *sigwait()* function can be implemented as a  
42800 macro on *sigwaitinfo()*.

42801 The *sigtimedwait()* function was separated from the *sigwaitinfo()* function to address concerns  
42802 regarding the overloading of the *timeout* pointer to indicate indefinite wait (no timeout), timed  
42803 wait, and immediate return, and concerns regarding consistency with other functions where the  
42804 conditional and timed waits were separate functions from the pure blocking function. The  
42805 semantics of *sigtimedwait()* are specified such that *sigwaitinfo()* could be implemented as a

42806 macro with a NULL pointer for *timeout*.

42807 The *sigwait* functions provide a synchronous mechanism for threads to wait for asynchronously  
 42808 generated signals. One important question was how many threads that are suspended in a call to  
 42809 a *sigwait()* function for a signal should return from the call when the signal is sent. Four choices  
 42810 were considered:

- 42811 1. Return an error for multiple simultaneous calls to *sigwait* functions for the same signal.
- 42812 2. One or more threads return.
- 42813 3. All waiting threads return.
- 42814 4. Exactly one thread returns.

42815 Prohibiting multiple calls to *sigwait()* for the same signal was felt to be overly restrictive. The  
 42816 “one or more” behavior made implementation of conforming packages easy at the expense of  
 42817 forcing POSIX threads clients to protect against multiple simultaneous calls to *sigwait()* in  
 42818 application code in order to achieve predictable behavior. There was concern that the “all  
 42819 waiting threads” behavior would result in “signal broadcast storms”, consuming excessive CPU  
 42820 resources by replicating the signals in the general case. Furthermore, no convincing examples  
 42821 could be presented that delivery to all was either simpler or more powerful than delivery to one.

42822 Thus, the consensus was that exactly one thread that was suspended in a call to a *sigwait*  
 42823 function for a signal should return when that signal occurs. This is not an onerous restriction as:

- 42824 • A multi-way signal wait can be built from the single-way wait.
- 42825 • Signals should only be handled by application-level code, as library routines cannot guess  
 42826 what the application wants to do with signals generated for the entire process.
- 42827 • Applications can thus arrange for a single thread to wait for any given signal and call any  
 42828 needed routines upon its arrival.

42829 In an application that is using signals for XSI interprocess communication, signal processing is  
 42830 typically done in one place. Alternatively, if the signal is being caught so that process cleanup  
 42831 can be done, the signal handler thread can call separate process cleanup routines for each  
 42832 portion of the application. Since the application main line started each portion of the application,  
 42833 it is at the right abstraction level to tell each portion of the application to clean up.

42834 Certainly, there exist programming styles where it is logical to consider waiting for a single  
 42835 signal in multiple threads. A simple *sigwait\_multiple()* routine can be constructed to achieve this  
 42836 goal. A possible implementation would be to have each *sigwait\_multiple()* caller registered as  
 42837 having expressed interest in a set of signals. The caller then waits on a thread-specific condition  
 42838 variable. A single server thread calls a *sigwait()* function on the union of all registered signals.  
 42839 When the *sigwait()* function returns, the appropriate state is set and condition variables are  
 42840 broadcast. New *sigwait\_multiple()* callers may cause the pending *sigwait()* call to be canceled  
 42841 and reissued in order to update the set of signals being waited for.

#### 42842 FUTURE DIRECTIONS

42843 None.

#### 42844 SEE ALSO

42845 Section 2.8.1 (on page 543), *pause()*, *pthread\_sigmask()*, *sigaction()*, *sigpending()*, *sigsuspend()*,  
 42846 *sigwait()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <signal.h>, <time.h>

42847 **CHANGE HISTORY**

42848 First released in Issue 5. Included for alignment with the POSIX Realtime Extension and the  
42849 POSIX Threads Extension.

42850 **Issue 6**

42851 These functions are marked as part of the Realtime Signals Extension option. |

42852 The Open Group corrigenda item U035/3 has been applied. The SYNOPSIS of the *sigwaitinfo()*  
42853 function has been corrected so that the second argument is of type **siginfo\_t\***.

42854 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
42855 implementation does not support the Realtime Signals Extension option. |

42856 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that the  
42857 CLOCK\_MONOTONIC clock, if supported, is used to measure timeout intervals. |

42858 The **restrict** keyword is added to the *sigtimedwait()* and *sigwaitinfo()* prototypes for alignment  
42859 with the ISO/IEC 9899:1999 standard. |

42860 **NAME**

42861 sigwait — wait for queued signals

42862 **SYNOPSIS**

42863 #include &lt;signal.h&gt;

42864 int sigwait(const sigset\_t \*restrict set, int \*restrict sig);

42865 **DESCRIPTION**

42866 The *sigwait()* function selects a pending signal from *set*, atomically clears it from the system's set  
 42867 of pending signals, and returns that signal number in the location referenced by *sig*. If prior to  
 42868 the call to *sigwait()* there are multiple pending instances of a single signal number, it is  
 42869 implementation-defined whether upon successful return there are any remaining pending  
 42870 RTS signals for that signal number. If the implementation supports queued signals and there are  
 42871 multiple signals queued for the signal number selected, the first such queued signal shall cause a  
 42872 return from *sigwait()* and the remainder shall remain queued. If no signal in *set* is pending at the  
 42873 time of the call, the thread is suspended until one or more becomes pending. The signals defined  
 42874 by *set* shall have been blocked at the time of the call to *sigwait()*; otherwise, the behavior is  
 42875 undefined. The effect of *sigwait()* on the signal actions for the signals in *set* is unspecified.

42876 If more than one thread is using *sigwait()* to wait for the same signal, no more than one of these  
 42877 threads shall return from *sigwait()* with the signal number. Which thread returns from *sigwait()*  
 42878 if more than a single thread is waiting is unspecified.

42879 RTS Should any of the multiple pending signals in the range SIGRTMIN to SIGRTMAX be selected, it  
 42880 shall be the lowest numbered one. The selection order between realtime and non-realtime  
 42881 signals, or between multiple pending non-realtime signals, is unspecified.

42882 **RETURN VALUE**

42883 Upon successful completion, *sigwait()* shall store the signal number of the received signal at the  
 42884 location referenced by *sig* and return zero. Otherwise, an error number shall be returned to  
 42885 indicate the error.

42886 **ERRORS**42887 The *sigwait()* function may fail if:42888 [EINVAL] The *set* argument contains an invalid or unsupported signal number.42889 **EXAMPLES**

42890 None.

42891 **APPLICATION USAGE**

42892 None.

42893 **RATIONALE**

42894 To provide a convenient way for a thread to wait for a signal, this volume of  
 42895 IEEE Std. 1003.1-200x provides the *sigwait()* function. For most cases where a thread has to wait  
 42896 for a signal, the *sigwait()* function should be quite convenient, efficient, and adequate.

42897 However, requests were made for a lower-level primitive than *sigwait()* and for semaphores that  
 42898 could be used by threads. After some consideration, threads were allowed to use semaphores  
 42899 and *sem\_post()* was defined to be async-signal and async-cancel-safe.

42900 In summary, when it is necessary for code run in response to an asynchronous signal to notify a  
 42901 thread, *sigwait()* should be used to handle the signal. Alternatively, if the implementation  
 42902 provides semaphores, they also can be used, either following *sigwait()* or from within a signal  
 42903 handling routine previously registered with *sigaction()*.

42904 **FUTURE DIRECTIONS**

42905           None.

42906 **SEE ALSO**

42907           Section 2.4 (on page 528), Section 2.8.1 (on page 543), *pause()*, *pthread\_sigmask()*, *sigaction()*,  
42908           *sigpending()*, *sigsuspend()*, *sigwaitinfo()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
42909           <**signal.h**>, <**time.h**>

42910 **CHANGE HISTORY**

42911           First released in Issue 5. Included for alignment with the POSIX Realtime Extension and the  
42912           POSIX Threads Extension.

42913 **Issue 6**

42914           The RATIONALE section is added.

42915           The **restrict** keyword is added to the *sigwait()* prototype for alignment with the  
42916           ISO/IEC 9899:1999 standard.

42917 **NAME**42918 sigwaitinfo — wait for queued signals (**REALTIME**)42919 **SYNOPSIS**

42920 RTS #include &lt;signal.h&gt;

42921 int sigwaitinfo(const sigset\_t \*restrict set, siginfo\_t \*restrict info); |

42922

42923 **DESCRIPTION**42924 Refer to *sigtimedwait()*.

42925 **NAME**

42926 sin, sinf, sinl — sine function

42927 **SYNOPSIS**

42928 #include &lt;math.h&gt;

42929 double sin(double x);

42930 float sinf(float x);

42931 long double sinl(long double x);

42932 **DESCRIPTION**

42933 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 42934 conflict between the requirements described here and the ISO C standard is unintentional. This  
 42935 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

42936 These functions shall compute the sine of its argument *x*, measured in radians.

42937 An application wishing to check for error situations should set *errno* to 0 before calling *sin()*. If  
 42938 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

42939 The *sin()* function may lose accuracy when its argument is far from 0.0.42940 **RETURN VALUE**42941 Upon successful completion, these functions shall return the sine of *x*.42942 **XSI** If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

42943 **XSI** If *x* is  $\pm\text{Inf}$ , either 0.0 shall be returned and *errno* set to [EDOM], or NaN shall be returned and  
 42944 *errno* may be set to [EDOM].

42945 If the correct result would cause underflow, 0.0 shall be returned and *errno* may be set to  
 42946 [ERANGE].

42947 **ERRORS**

42948 These functions may fail if:

42949 **XSI** [EDOM] The value of *x* is NaN, or *x* is  $\pm\text{Inf}$ .

42950 [ERANGE] The result underflows.

42951 **XSI** No other errors shall occur.42952 **EXAMPLES**42953 **Taking the Sine of a 45-Degree Angle**

42954 #include &lt;math.h&gt;

42955 ...

42956 double radians = 45.0 \* M\_PI / 180;

42957 double result;

42958 ...

42959 result = sin(radians);

42960 **APPLICATION USAGE**

42961 None.

42962 **RATIONALE**

42963 None.

42964 **FUTURE DIRECTIONS**

42965 None.

42966 **SEE ALSO**42967 *asin()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>42968 **CHANGE HISTORY**

42969 First released in Issue 1. Derived from Issue 1 of the SVID.

42970 **Issue 4**42971 References to *matherr()* are removed.42972 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
42973 ISO C standard and to rationalize error handling in the mathematics functions.

42974 The return value specified for [EDOM] is marked as an extension.

42975 **Issue 5**42976 The last two paragraphs of the DESCRIPTION were included as APPLICATION USAGE notes  
42977 in previous issues.42978 **Issue 6**42979 The *sinf()* and *sinl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

42980 **NAME**

42981           sinh, sinhf, sinhl — hyperbolic sine function

42982 **SYNOPSIS**

42983           #include <math.h>

42984           double sinh(double x);

42985           float sinhf(float x);

42986           long double sinhl(long double x);

42987 **DESCRIPTION**

42988 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
42989 conflict between the requirements described here and the ISO C standard is unintentional. This  
42990 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

42991           These functions shall compute the hyperbolic sine of  $x$ .

42992           An application wishing to check for error situations should set *errno* to 0 before calling *sinh()*. If  
42993 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

42994 **RETURN VALUE**

42995           Upon successful completion, these functions shall return the hyperbolic sine of  $x$ .

42996           If the result would cause an overflow,  $\pm$ HUGE\_VAL shall be returned and *errno* set to  
42997 [ERANGE].

42998           If the result would cause underflow, 0.0 shall be returned and *errno* may be set to [ERANGE].

42999 **XSI**       If  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].

43000 **ERRORS**

43001           These functions shall fail if:

43002           [ERANGE]       The result would cause overflow.

43003           These functions may fail if:

43004 **XSI**       [EDOM]       The value of  $x$  is NaN.

43005           [ERANGE]       The result would cause underflow.

43006 **XSI**       No other errors shall occur.

43007 **EXAMPLES**

43008           None.

43009 **APPLICATION USAGE**

43010           None.

43011 **RATIONALE**

43012           None.

43013 **FUTURE DIRECTIONS**

43014           None.

43015 **SEE ALSO**

43016           *asinh()*, *cosh()*, *isnan()*, *tanh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

43017 **CHANGE HISTORY**

43018           First released in Issue 1. Derived from Issue 1 of the SVID.

43019 **Issue 4**

43020 References to *matherr()* are removed.

43021 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
43022 ISO C standard and to rationalize error handling in the mathematics functions.

43023 The return value specified for [EDOM] is marked as an extension.

43024 **Issue 5**

43025 The DESCRIPTION is updated to indicate how an application should check for an error. This  
43026 text was previously published in the APPLICATION USAGE section.

43027 **Issue 6**

43028 The *sinhf()* and *sinhl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

43029 **NAME**

43030 sleep — suspend execution for an interval of time

43031 **SYNOPSIS**

43032 #include <unistd.h>

43033 unsigned sleep(unsigned *seconds*);

43034 **DESCRIPTION**

43035 The *sleep()* function shall cause the calling thread to be suspended from execution until either  
43036 the number of realtime seconds specified by the argument *seconds* has elapsed or a signal is  
43037 delivered to the calling thread and its action is to invoke a signal-catching function or to  
43038 terminate the process. The suspension time may be longer than requested due to the scheduling  
43039 of other activity by the system.

43040 If a SIGALRM signal is generated for the calling process during execution of *sleep()* and if the  
43041 SIGALRM signal is being ignored or blocked from delivery, it is unspecified whether *sleep()*  
43042 returns when the SIGALRM signal is scheduled. If the signal is being blocked, it is also  
43043 unspecified whether it remains pending after *sleep()* returns or it is discarded.

43044 If a SIGALRM signal is generated for the calling process during execution of *sleep()*, except as a  
43045 result of a prior call to *alarm()*, and if the SIGALRM signal is not being ignored or blocked from  
43046 delivery, it is unspecified whether that signal has any effect other than causing *sleep()* to return.

43047 If a signal-catching function interrupts *sleep()* and examines or changes either the time a  
43048 SIGALRM is scheduled to be generated, the action associated with the SIGALRM signal, or  
43049 whether the SIGALRM signal is blocked from delivery, the results are unspecified.

43050 If a signal-catching function interrupts *sleep()* and calls *siglongjmp()* or *longjmp()* to restore an  
43051 environment saved prior to the *sleep()* call, the action associated with the SIGALRM signal and  
43052 the time at which a SIGALRM signal is scheduled to be generated are unspecified. It is also  
43053 unspecified whether the SIGALRM signal is blocked, unless the process' signal mask is restored  
43054 as part of the environment.

43055 XSI Interactions between *sleep()* and any of *setitimer()*, *ualarm()*, or *usleep()* are unspecified.

43056 **RETURN VALUE**

43057 If *sleep()* returns because the requested time has elapsed, the value returned shall be 0. If *sleep()*  
43058 returns because of premature arousal due to delivery of a signal, the return value shall be the  
43059 “unslept” amount (the requested time minus the time actually slept) in seconds.

43060 **ERRORS**

43061 No errors are defined.

43062 **EXAMPLES**

43063 None.

43064 **APPLICATION USAGE**

43065 None.

43066 **RATIONALE**

43067 There are two general approaches to the implementation of the *sleep()* function. One is to use the  
43068 *alarm()* function to schedule a SIGALRM signal and then suspend the process waiting for that  
43069 signal. The other is to implement an independent facility. This volume of IEEE Std. 1003.1-200x  
43070 permits either approach.

43071 In order to comply with the requirement that no primitive shall change a process attribute unless  
43072 explicitly described by this volume of IEEE Std. 1003.1-200x, an implementation using SIGALRM  
43073 must carefully take into account any SIGALRM signal scheduled by previous *alarm()* calls, the

43074 action previously established for SIGALRM, and whether SIGALRM was blocked. If a SIGALRM  
 43075 has been scheduled before the *sleep()* would ordinarily complete, the *sleep()* must be shortened  
 43076 to that time and a SIGALRM generated (possibly simulated by direct invocation of the signal-  
 43077 catching function) before *sleep()* returns. If a SIGALRM has been scheduled after the *sleep()*  
 43078 would ordinarily complete, it must be rescheduled for the same time before *sleep()* returns. The  
 43079 action and blocking for SIGALRM must be saved and restored.

43080 Historical implementations often implement the SIGALRM-based version using *alarm()* and  
 43081 *pause()*. One such implementation is prone to infinite hangups, as described in *pause()*. Another  
 43082 such implementation uses the C-language *setjmp()* and *longjmp()* functions to avoid that  
 43083 window. That implementation introduces a different problem: when the SIGALRM signal  
 43084 interrupts a signal-catching function installed by the user to catch a different signal, the  
 43085 *longjmp()* aborts that signal-catching function. An implementation based on *sigprocmask()*,  
 43086 *alarm()*, and *sigsuspend()* can avoid these problems.

43087 Despite all reasonable care, there are several very subtle, but detectable and unavoidable,  
 43088 differences between the two types of implementations. These are the cases mentioned in this  
 43089 volume of IEEE Std. 1003.1-200x where some other activity relating to SIGALRM takes place,  
 43090 and the results are stated to be unspecified. All of these cases are sufficiently unusual as not to  
 43091 be of concern to most applications.

43092 See also the discussion of the term *realtime* in *alarm()*.

43093 Because *sleep()* can be implemented using *alarm()*, the discussion about alarms occurring early  
 43094 under *alarm()* applies to *sleep()* as well.

43095 Application writers should note that the type of the argument *seconds* and the return value of  
 43096 *sleep()* is **unsigned**. That means that a Strictly Conforming POSIX System Interfaces Application  
 43097 cannot pass a value greater than the minimum guaranteed value for {UINT\_MAX}, which the  
 43098 ISO C standard sets as 65 535, and any application passing a larger value is restricting its  
 43099 portability. A different type was considered, but historical implementations, including those  
 43100 with a 16-bit **int** type, consistently use either **unsigned** or **int**.

43101 Scheduling delays may cause the process to return from the *sleep()* function significantly after  
 43102 the requested time. In such cases, the return value should be set to zero, since the formula  
 43103 (requested time minus the time actually spent) yields a negative number and *sleep()* returns an  
 43104 **unsigned**.

#### 43105 FUTURE DIRECTIONS

43106 None.

#### 43107 SEE ALSO

43108 *alarm()*, *getitimer()*, *nanosleep()*, *pause()*, *sigaction()*, *sigsetjmp()*, *ualarm()*, *usleep()*, the Base  
 43109 Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

#### 43110 CHANGE HISTORY

43111 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 43112 Issue 4

43113 The <**unistd.h**> header is added to the SYNOPSIS section.

#### 43114 Issue 4, Version 2

43115 The DESCRIPTION is updated to indicate possible interactions with the *setitimer()*, *ualarm()*,  
 43116 and *usleep()* functions.

43117 **Issue 5**

43118

The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

43119 **NAME**

43120 socket — create an endpoint for communication

43121 **SYNOPSIS**

43122 #include &lt;sys/socket.h&gt;

43123 int socket(int *domain*, int *type*, int *protocol*);43124 **DESCRIPTION**43125 The *socket()* function shall create an unbound socket in a communications domain, and return a  
43126 file descriptor that can be used in later function calls that operate on sockets.43127 The *socket()* function takes the following arguments:43128 *domain* Specifies the communications domain in which a socket is to be created.43129 *type* Specifies the type of socket to be created.43130 *protocol* Specifies a particular protocol to be used with the socket. Specifying a *protocol*  
43131 of 0 causes *socket()* to use an unspecified default protocol appropriate for the  
43132 requested socket type.43133 The *domain* argument specifies the address family used in the communications domain. The  
43134 address families supported by the system are implementation-defined.43135 Symbolic constants that can be used for the domain argument are defined in the <sys/socket.h>  
43136 header.43137 The *type* argument specifies the socket type, which determines the semantics of communication  
43138 over the socket. The socket types supported by the system are implementation-defined. Possible  
43139 socket types include:43140 SOCK\_STREAM Provides sequenced, reliable, bidirectional, connection-mode byte  
43141 streams, and may provide a transmission mechanism for out-of-band  
43142 data.43143 SOCK\_DGRAM Provides datagrams, which are connectionless-mode, unreliable messages  
43144 of fixed maximum length.43145 SOCK\_SEQPACKET Provides sequenced, reliable, bidirectional, connection-mode  
43146 transmission path for records. A record can be sent using one or more  
43147 output operations and received using one or more input operations, but a  
43148 single operation never transfers part of more than one record. Record  
43149 boundaries are visible to the receiver via the MSG\_EOR flag.43150 If the *protocol* argument is non-zero, it shall specify a protocol that is supported by the address  
43151 family. The protocols supported by the system are implementation-defined.43152 The process may need to have appropriate privileges to use the *socket()* function or to create  
43153 some sockets.43154 **RETURN VALUE**43155 Upon successful completion, *socket()* shall return a non-negative integer, the socket file  
43156 descriptor. Otherwise, a value of -1 shall be returned and *errno* set to indicate the error.43157 **ERRORS**43158 The *socket()* function shall fail if:

43159 [EAFNOSUPPORT]

43160 The implementation does not support the specified address family.

- 43161 [EMFILE] No more file descriptors are available for this process.
- 43162 [ENFILE] No more file descriptors are available for the system.
- 43163 [EPROTONOSUPPORT]  
43164 The protocol is not supported by the address family, or the protocol is not  
43165 supported by the implementation.
- 43166 [EPROTOYPE] The socket type is not supported by the protocol.
- 43167 The *socket()* function may fail if:
- 43168 [EACCES] The process does not have appropriate privileges.
- 43169 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 43170 [ENOMEM] Insufficient memory was available to fulfill the request.
- 43171 **EXAMPLES**
- 43172 None.
- 43173 **APPLICATION USAGE**
- 43174 The documentation for specific address families specifies which protocols each address family  
43175 supports. The documentation for specific protocols specifies which socket types each protocol  
43176 supports.
- 43177 The application can determine if an address family is supported by trying to create a socket with  
43178 *domain* set to the protocol in question.
- 43179 **RATIONALE**
- 43180 None.
- 43181 **FUTURE DIRECTIONS**
- 43182 None.
- 43183 **SEE ALSO**
- 43184 *accept()*, *bind()*, *connect()*, *getsockname()*, *getsockopt()*, *listen()*, *recv()*, *recvfrom()*, *recvmsg()*,  
43185 *send()*, *sendmsg()*, *setsockopt()*, *shutdown()*, *socketpair()*, the Base Definitions volume of  
43186 IEEE Std. 1003.1-200x, <netinet/in.h>, <sys/socket.h>
- 43187 **CHANGE HISTORY**
- 43188 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

43189 **NAME**

43190 socketpair — create a pair of connected sockets

43191 **SYNOPSIS**

43192 #include &lt;sys/socket.h&gt;

43193 int socketpair(int *domain*, int *type*, int *protocol*,  
43194 int *socket\_vector*[2]);43195 **DESCRIPTION**43196 The *socketpair()* function creates an unbound pair of connected sockets in a specified *domain*, of a  
43197 specified *type*, under the protocol optionally specified by the *protocol* argument. The two sockets  
43198 are identical. The file descriptors used in referencing the created sockets are returned in  
43199 *socket\_vector*[0] and *socket\_vector*[1].43200 The *socketpair()* function takes the following arguments:

|       |                      |                                                                                                                                                                                                                               |
|-------|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 43201 | <i>domain</i>        | Specifies the communications domain in which the sockets are to be created.                                                                                                                                                   |
| 43202 | <i>type</i>          | Specifies the type of sockets to be created.                                                                                                                                                                                  |
| 43203 | <i>protocol</i>      | Specifies a particular protocol to be used with the sockets. Specifying a<br>43204 <i>protocol</i> of 0 causes <i>socketpair()</i> to use an unspecified default protocol<br>43205 appropriate for the requested socket type. |
| 43206 | <i>socket_vector</i> | Specifies a 2-integer array to hold the file descriptors of the created socket<br>43207 pair.                                                                                                                                 |

43208 The *type* argument specifies the socket type, which determines the semantics of communications  
43209 over the socket. The socket types supported by the system are implementation-defined. Possible  
43210 socket types include:

|       |                |                                                                                                                                                                                                                                                                                                                                                                                     |
|-------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 43211 | SOCK_STREAM    | Provides sequenced, reliable, bidirectional, connection-mode byte<br>43212 streams, and may provide a transmission mechanism for out-of-band<br>43213 data.                                                                                                                                                                                                                         |
| 43214 | SOCK_DGRAM     | Provides datagrams, which are connectionless-mode, unreliable messages<br>43215 of fixed maximum length.                                                                                                                                                                                                                                                                            |
| 43216 | SOCK_SEQPACKET | Provides sequenced, reliable, bidirectional, connection-mode<br>43217 transmission paths for records. A record can be sent using one or more<br>43218 output operations and received using one or more input operations, but a<br>43219 single operation never transfers part of more than one record. Record<br>43220 boundaries are visible to the receiver via the MSG_EOR flag. |

43221 If the *protocol* argument is non-zero, it shall specify a protocol that is supported by the address  
43222 family. The protocols supported by the system are implementation-defined.43223 The process may need to have appropriate privileges to use the *socketpair()* function or to create  
43224 some sockets.43225 **RETURN VALUE**43226 Upon successful completion, this function shall return 0; otherwise, -1 shall be returned and  
43227 *errno* set to indicate the error.43228 **ERRORS**43229 The *socketpair()* function shall fail if:

43230 [EAFNOSUPPORT]

43231 The implementation does not support the specified address family.

- 43232 [EMFILE] No more file descriptors are available for this process.
- 43233 [ENFILE] No more file descriptors are available for the system.
- 43234 [EOPNOTSUPP] The specified protocol does not permit creation of socket pairs.
- 43235 [EPROTONOSUPPORT]  
43236 The protocol is not supported by the address family, or the protocol is not  
43237 supported by the implementation.
- 43238 [EPROTOTYPE] The socket type is not supported by the protocol.
- 43239 The *socketpair()* function may fail if:
- 43240 [EACCES] The process does not have appropriate privileges.
- 43241 [ENOBUFS] Insufficient resources were available in the system to perform the operation.
- 43242 [ENOMEM] Insufficient memory was available to fulfill the request.
- 43243 **EXAMPLES**
- 43244 None.
- 43245 **APPLICATION USAGE**
- 43246 The documentation for specific address families specifies which protocols each address family  
43247 supports. The documentation for specific protocols specifies which socket types each protocol  
43248 supports.
- 43249 The *socketpair()* function is used primarily with UNIX domain sockets and need not be  
43250 supported for other domains.
- 43251 **RATIONALE**
- 43252 None.
- 43253 **FUTURE DIRECTIONS**
- 43254 None.
- 43255 **SEE ALSO**
- 43256 *socket()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/socket.h>
- 43257 **CHANGE HISTORY**
- 43258 First released in Issue 6. Derived from the XNS, Issue 5.2 specification.

43259 **NAME**

43260        sprintf, snprintf — print formatted output

43261 **SYNOPSIS**

43262        #include &lt;stdio.h&gt;

43263        int snprintf(char \*restrict *s*, size\_t *n*,43264            const char \*restrict *format*, /\* args \*/ ...);43265        int sprintf(char \*restrict *s*, const char \*restrict *format*, ...);43266 **DESCRIPTION**43267        Refer to *fprintf()*.

43268 **NAME**

43269 sqrt, sqrtf, sqrtl — square root function

43270 **SYNOPSIS**

43271 #include &lt;math.h&gt;

43272 double sqrt(double x);

43273 float sqrtf(float x);

43274 long double sqrtl(long double x);

43275 **DESCRIPTION**

43276 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 43277 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43278 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

43279 These functions shall compute the square root of  $x$ ,  $\sqrt{x}$ .

43280 An application wishing to check for error situations should set *errno* to 0 before calling *sqrt()*. If  
 43281 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

43282 **RETURN VALUE**43283 Upon successful completion, these functions shall return the square root of  $x$ .43284 **XSI** If  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].43285 **XSI** If  $x$  is negative, 0.0 or NaN shall be returned and *errno* shall be set to [EDOM].43286 **ERRORS**

43287 These functions shall fail if:

43288 [EDOM] The value of  $x$  is negative.

43289 These functions may fail if:

43290 **XSI** [EDOM] The value of  $x$  is NaN.43291 **XSI** No other errors shall occur.43292 **EXAMPLES**43293 **Taking the Square Root of 9.0**

43294 #include &lt;math.h&gt;

43295 ...

43296 double x = 9.0;

43297 double result;

43298 ...

43299 result = sqrt(x);

43300 **APPLICATION USAGE**

43301 None.

43302 **RATIONALE**

43303 None.

43304 **FUTURE DIRECTIONS**

43305 None.

43306 **SEE ALSO**

43307 *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>, <stdio.h>

43308 **CHANGE HISTORY**

43309 First released in Issue 1. Derived from Issue 1 of the SVID.

43310 **Issue 4**

43311 References to *matherr()* are removed.

43312 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
43313 ISO C standard and to rationalize error handling in the mathematics functions.

43314 The return value specified for [EDOM] is marked as an extension.

43315 **Issue 5**

43316 The DESCRIPTION is updated to indicate how an application should check for an error. This  
43317 text was previously published in the APPLICATION USAGE section.

43318 **Issue 6**

43319 The *sqrtrf()* and *sqrtrl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

43320 **NAME**

43321        `srand` — pseudo-random number generator

43322 **SYNOPSIS**

43323        `#include <stdlib.h>`

43324        `void srand(unsigned seed);`

43325 **DESCRIPTION**

43326        Refer to `rand()`.

43327 **NAME**

43328           srand48 — seed uniformly distributed double-precision pseudo-random number generator

43329 **SYNOPSIS**

43330 XSI        #include &lt;stdlib.h&gt;

43331            void srand48(long *seedval*);

43332

43333 **DESCRIPTION**43334            Refer to *drand48()*.

43335 **NAME**

43336           srandom — seed pseudo-random number generator

43337 **SYNOPSIS**

43338 XSI       #include <stdlib.h>

43339           void srandom(unsigned *seed*);

43340

43341 **DESCRIPTION**

43342           Refer to *initstate()*.

43343 **NAME**

43344        sscanf — convert formatted input

43345 **SYNOPSIS**

43346        #include &lt;stdio.h&gt;

43347        int sscanf(const char \*restrict *s*, const char \*restrict *format*, ...); |43348 **DESCRIPTION** |43349        Refer to *fscanf()*.

## 43350 NAME

43351 stat — get file status

## 43352 SYNOPSIS

43353 #include &lt;sys/stat.h&gt;

43354 int stat(const char \*restrict path, struct stat \*restrict buf);

## 43355 DESCRIPTION

43356 The *stat()* function obtains information about the named file and writes it to the area pointed to  
 43357 by the *buf* argument. The *path* argument points to a path name naming a file. Read, write, or  
 43358 execute permission of the named file is not required. An implementation that provides  
 43359 additional or alternate file access control mechanisms may, under implementation-defined  
 43360 conditions, cause *stat()* to fail. In particular, the system may deny the existence of the file  
 43361 specified by *path*.

43362 If the named file is a symbolic link, the *stat()* function shall continue path name resolution using  
 43363 the contents of the symbolic link, and shall return information pertaining to the resulting file if  
 43364 the file exists.

43365 The *buf* argument is a pointer to a **stat** structure, as defined in the header <sys/stat.h>, into  
 43366 which information is placed concerning the file.

43367 The *stat()* function updates any time-related fields (as described in the definition of **File Times**  
 43368 **Update** in the Base Definitions volume of IEEE Std. 1003.1-200x), before writing into the **stat**  
 43369 structure.

43370 The structure members *st\_mode*, *st\_ino*, *st\_dev*, *st\_uid*, *st\_gid*, *st\_atime*, *st\_ctime*, and *st\_mtime*  
 43371 shall have meaningful values for all file types defined in this volume of IEEE Std. 1003.1-200x.  
 43372 The value of the member *st\_nlink* shall be set to the number of links to the file.

## 43373 RETURN VALUE

43374 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
 43375 indicate the error.

## 43376 ERRORS

43377 The *stat()* function shall fail if:

43378 [EACCES] Search permission is denied for a component of the path prefix.

43379 [EIO] An error occurred while reading from the file system.

43380 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
43381 argument.

43382 [ENAMETOOLONG]

43383 The length of the *path* argument exceeds {PATH\_MAX} or a path name  
43384 component is longer than {NAME\_MAX}.43385 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

43386 [ENOTDIR] A component of the path prefix is not a directory.

43387 [EOVERFLOW] The file size in bytes or the number of blocks allocated to the file or the file  
43388 serial number cannot be represented correctly in the structure pointed to by  
43389 *buf*.43390 The *stat()* function may fail if:43391 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
43392 resolution of the *path* argument.

43393 [ENAMETOOLONG]  
 43394 As a result of encountering a symbolic link in resolution of the *path* argument,  
 43395 the length of the substituted path name string exceeded {PATH\_MAX}.

43396 [EOVERFLOW] A value to be stored would overflow one of the members of the **stat** structure.

43397 **EXAMPLES**43398 **Obtaining File Status Information**

43399 The following example shows how to obtain file status information for a file named  
 43400 **/home/cnd/mod1**. The structure variable *buffer* is defined for the **stat** structure.

```
43401 #include <sys/types.h>
43402 #include <sys/stat.h>
43403 #include <fcntl.h>

43404 struct stat buffer;
43405 int status;
43406 ...
43407 status = stat("/home/cnd/mod1", &buffer);
```

43408 **Getting Directory Information**

43409 The following example fragment gets status information for each entry in a directory. The call to  
 43410 the *stat()* function stores file information in the **stat** structure pointed to by *statbuf*. The lines  
 43411 that follow the *stat()* call format the fields in the **stat** structure for presentation to the user of the  
 43412 program.

```
43413 #include <sys/types.h>
43414 #include <sys/stat.h>
43415 #include <dirent.h>
43416 #include <pwd.h>
43417 #include <grp.h>
43418 #include <time.h>
43419 #include <locale.h>
43420 #include <langinfo.h>

43421 struct dirent *dp;
43422 struct stat statbuf;
43423 struct passwd *pwd;
43424 struct group *grp;
43425 struct tm *tm;
43426 char datestring[256];
43427 ...
43428 /* Loop through directory entries */
43429 while ((dp = readdir(dir)) != NULL) {

43430 /* Get entry's information. */
43431 if (stat(dp->d_name, &statbuf) == -1)
43432 continue;

43433 /* Print out type, permissions, and number of links. */
43434 printf("%10.10s", sperm (statbuf.st_mode));
43435 printf("%4d", statbuf.st_nlink);
```

```

43436 /* Print out owners name if it is found using getpwuid(). */
43437 if ((pwd = getpwuid(statbuf.st_uid)) != NULL)
43438 printf(" %-8.8s", pwd->pw_name);
43439 else
43440 printf(" %-8d", statbuf.st_uid);
43441
43442 /* Print out group name if it's found using getgrgid(). */
43443 if ((grp = getgrgid(statbuf.st_gid)) != NULL)
43444 printf(" %-8.8s", grp->gr_name);
43445 else
43446 printf(" %-8d", statbuf.st_gid);
43447
43448 /* Print size of file. */
43449 printf("%9ld", statbuf.st_size);
43450
43451 tm = localtime(&statbuf.st_mtime);
43452
43453 /* Get localized date string. */
43454 strftime(datestring, sizeof(datestring), nl_langinfo(D_T_FMT), tm);
43455 printf(" %s %s\n", datestring, dp->d_name);
43456 }

```

#### 43453 APPLICATION USAGE

43454 None.

#### 43455 RATIONALE

43456 The intent of the paragraph describing “additional or alternate file access control mechanisms”  
 43457 is to allow a secure implementation where a process with a label that does not dominate the  
 43458 file’s label cannot perform a *stat()* function. This is not related to read permission; a process with  
 43459 a label that dominates the file’s label does not need read permission. An implementation that  
 43460 supports write-up operations could fail *fstat()* function calls even though it has a valid file  
 43461 descriptor open for writing.

#### 43462 FUTURE DIRECTIONS

43463 None.

#### 43464 SEE ALSO

43465 *fstat()*, *lstat()*, *readlink()*, *symlink()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
 43466 <sys/stat.h>, <sys/types.h>

#### 43467 CHANGE HISTORY

43468 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 43469 Issue 4

43470 The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
 43471 XSI-conformant systems.

43472 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 43473 • The type of argument *path* is changed from **char\*** to **const char\***.
- 43474 • In the DESCRIPTION is changed as follows:
  - 43475 — Statements indicating the purpose of this function and a paragraph defining the contents
  - 43476 of **stat** structure members are added.
  - 43477 — The words “extended security controls” are replaced by “additional or alternate file
  - 43478 access control mechanisms”.

- 43479 The following change is incorporated for alignment with the FIPS requirements:
- 43480 • In the **ERRORS** section, the condition whereby [ENAMETOOLONG] is returned if a path
  - 43481 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as
  - 43482 an extension.
- 43483 **Issue 4, Version 2**
- 43484 The **ERRORS** section is updated for X/OPEN UNIX conformance as follows:
- 43485 • In the mandatory section, [EIO] is added to indicate that a physical I/O error has occurred,
  - 43486 and [ELOOP] to indicate that too many symbolic links were encountered during path name
  - 43487 resolution.
  - 43488 • In the optional section, a second [ENAMETOOLONG] condition is defined that may report
  - 43489 excessive length of an intermediate result of path name resolution of a symbolic link.
  - 43490 • In the optional section, [EOVERFLOW] is added to indicate that a value to be stored in a
  - 43491 member of the **stat** structure would cause overflow.
- 43492 **Issue 5**
- 43493 Large File Summit extensions are added.
- 43494 **Issue 6**
- 43495 In the **SYNOPSIS**, the inclusion of `<sys/types.h>` is no longer required.
- 43496 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:
- 43497 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.
  - 43498 This is since behavior may vary from one file system to another.
- 43499 The following new requirements on POSIX implementations derive from alignment with the
- 43500 Single UNIX Specification:
- 43501 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was
  - 43502 required for conforming implementations of previous POSIX specifications, it was not
  - 43503 required for UNIX applications.
  - 43504 • The [EIO] mandatory error condition is added.
  - 43505 • The [ELOOP] mandatory error condition is added.
  - 43506 • The [EOVERFLOW] mandatory error condition is added. This change is to support large
  - 43507 files.
  - 43508 • The [ENAMETOOLONG] and the second [EOVERFLOW] optional error conditions are
  - 43509 added.
- 43510 The following changes were made to align with the IEEE P1003.1a draft standard:
- 43511 • Details are added regarding the treatment of symbolic links.
  - 43512 • The [ELOOP] optional error condition is added.
- 43513 The **DESCRIPTION** is updated to avoid use of the term “must” for application requirements.
- 43514 The **restrict** keyword is added to the `stat()` prototype for alignment with the ISO/IEC 9899: 1999
- 43515 standard.

43516 **NAME**

43517           statvfs — get file system information

43518 **SYNOPSIS**

43519 XSI       #include <sys/statvfs.h>

43520           int statvfs(const char \*restrict path, struct statvfs \*restrict buf);

43521

43522 **DESCRIPTION**

43523           Refer to *fstatvfs()*.

43524 **NAME**

43525 stderr, stdin, stdout — standard I/O streams

43526 **SYNOPSIS**

43527 #include &lt;stdio.h&gt;

43528 extern FILE \**stderr*, \**stdin*, \**stdout*;43529 **DESCRIPTION**

43530 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 43531 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43532 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

43533 A file with associated buffering is called a *stream* and is declared to be a pointer to a defined type  
 43534 **FILE**. The *fopen()* function creates certain descriptive data for a stream and returns a pointer to  
 43535 designate the stream in all further transactions. Normally, there are three open streams with  
 43536 constant pointers declared in the <stdio.h> header and associated with the standard open files.

43537 At program start-up, three streams are predefined and need not be opened explicitly: *standard*  
 43538 *input* (for reading conventional input), *standard output* (for writing conventional output), and  
 43539 *standard error* (for writing diagnostic output). When opened, the standard error stream is not  
 43540 fully buffered; the standard input and standard output streams are fully buffered if and only if  
 43541 the stream can be determined not to refer to an interactive device.

43542 CX The following symbolic values in <unistd.h> define the file descriptors that shall be associated  
 43543 with the C-language *stdin*, *stdout*, and *stderr* when the application is started:

43544 STDIN\_FILENO Standard input value, *stdin*. Its value is 0.

43545 STDOUT\_FILENO Standard output value, *stdout*. Its value is 1.

43546 STDERR\_FILENO Standard error value, *stderr*. Its value is 2.

43547

43548 *stderr* is expected to be open for reading and writing.

43549 **RETURN VALUE**

43550 None.

43551 **ERRORS**

43552 No errors are defined.

43553 **EXAMPLES**

43554 None.

43555 **APPLICATION USAGE**

43556 None.

43557 **RATIONALE**

43558 None.

43559 **FUTURE DIRECTIONS**

43560 None.

43561 **SEE ALSO**

43562 *fclose()*, *feof()*, *ferror()*, *fileno()*, *fopen()*, *fread()*, *fseek()*, *getc()*, *gets()*, *popen()*, *printf()*, *putc()*,  
 43563 *puts()*, *read()*, *scanf()*, *setbuf()*, *setvbuf()*, *tmpfile()*, *ungetc()*, *vprintf()*, the Base Definitions  
 43564 volume of IEEE Std. 1003.1-200x, <stdio.h>, <unistd.h>

43565 **CHANGE HISTORY**

43566           First released in Issue 1.

43567 **Issue 6**

43568           Extensions beyond the ISO C standard are now marked.

43569 **NAME**43570 `strcasecmp`, `strncasecmp` — case-insensitive string comparisons43571 **SYNOPSIS**43572 XSI `#include <strings.h>`43573 `int strcmp(const char *s1, const char *s2);`43574 `int strncasecmp(const char *s1, const char *s2, size_t n);`

43575

43576 **DESCRIPTION**

43577 The `strcasecmp()` function compares, while ignoring differences in case, the string pointed to by  
43578 `s1` to the string pointed to by `s2`. The `strncasecmp()` function compares, while ignoring  
43579 differences in case, not more than `n` bytes from the string pointed to by `s1` to the string pointed to  
43580 by `s2`.

43581 In the POSIX locale, `strcasecmp()` and `strncasecmp()` do upper to lower conversions, then a byte  
43582 comparison. The results are unspecified in other locales.

43583 **RETURN VALUE**

43584 Upon completion, `strcasecmp()` shall return an integer greater than, equal to, or less than 0, if the  
43585 string pointed to by `s1` is, ignoring case, greater than, equal to, or less than the string pointed to  
43586 by `s2`, respectively.

43587 Upon successful completion, `strncasecmp()` shall return an integer greater than, equal to, or less  
43588 than 0, if the possibly null-terminated array pointed to by `s1` is, ignoring case, greater than, equal  
43589 to, or less than the possibly null-terminated array pointed to by `s2`, respectively.

43590 **ERRORS**

43591 No errors are defined.

43592 **EXAMPLES**

43593 None.

43594 **APPLICATION USAGE**

43595 None.

43596 **RATIONALE**

43597 None.

43598 **FUTURE DIRECTIONS**

43599 None.

43600 **SEE ALSO**43601 The Base Definitions volume of IEEE Std. 1003.1-200x, `<strings.h>`43602 **CHANGE HISTORY**

43603 First released in Issue 4, Version 2.

43604 **Issue 5**

43605 Moved from X/OPEN UNIX extension to BASE.

43606 **NAME**

43607           strcat — concatenate two strings

43608 **SYNOPSIS**

43609           #include &lt;string.h&gt;

43610           char \*strcat(char \*restrict *s1*, const char \*restrict *s2*);43611 **DESCRIPTION**

43612 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
43613 conflict between the requirements described here and the ISO C standard is unintentional. This  
43614 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

43615       The *strcat()* function shall append a copy of the string pointed to by *s2* (including the  
43616 terminating null byte) to the end of the string pointed to by *s1*. The initial byte of *s2* overwrites  
43617 the null byte at the end of *s1*. If copying takes place between objects that overlap, the behavior is  
43618 undefined.

43619 **RETURN VALUE**43620       The *strcat()* function shall return *s1*; no return value is reserved to indicate an error.43621 **ERRORS**

43622       No errors are defined.

43623 **EXAMPLES**

43624       None.

43625 **APPLICATION USAGE**

43626       This issue is aligned with the ISO C standard; this does not affect compatibility with XPG3  
43627 applications. Reliable error detection by this function was never guaranteed.

43628 **RATIONALE**

43629       None.

43630 **FUTURE DIRECTIONS**

43631       None.

43632 **SEE ALSO**43633       *strncat()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>43634 **CHANGE HISTORY**

43635       First released in Issue 1. Derived from Issue 1 of the SVID.

43636 **Issue 4**

43637       The **DESCRIPTION** is changed to make it clear that the function manipulates bytes rather than  
43638 (possibly multi-byte) characters.

43639       The following change is incorporated for alignment with the ISO C standard:

- 43640
  - The type of argument *s2* is changed from **char\*** to **const char\***.

43641 **Issue 6**43642       The *strcat()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

43643 **NAME**

43644           strchr — string scanning operation

43645 **SYNOPSIS**

43646           #include &lt;string.h&gt;

43647           char \*strchr(const char \*s, int c);

43648 **DESCRIPTION**43649 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
43650 conflict between the requirements described here and the ISO C standard is unintentional. This  
43651 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.43652 **CX**       The *strchr()* function shall locate the first occurrence of *c* (converted to an **unsigned char**) in the  
43653 string pointed to by *s*. The terminating null byte is considered to be part of the string.43654 **RETURN VALUE**43655           Upon completion, *strchr()* shall return a pointer to the byte, or a null pointer if the byte was not  
43656 found.43657 **ERRORS**

43658           No errors are defined.

43659 **EXAMPLES**

43660           None.

43661 **APPLICATION USAGE**

43662           None.

43663 **RATIONALE**

43664           None.

43665 **FUTURE DIRECTIONS**

43666           None.

43667 **SEE ALSO**43668           *strchr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>43669 **CHANGE HISTORY**

43670           First released in Issue 1. Derived from Issue 1 of the SVID.

43671 **Issue 4**43672           The **DESCRIPTION** and **RETURN VALUE** sections are changed to make it clear that the function  
43673 manipulates bytes rather than (possibly multi-byte) characters.43674           The **APPLICATION USAGE** section is removed.

43675           The following change is incorporated for alignment with the ISO C standard:

- 43676
- The type of argument *s* is changed from **char\*** to **const char\***.

43677 **Issue 6**

43678           Extensions beyond the ISO C standard are now marked.

43679 **NAME**43680 `stricmp` — compare two strings43681 **SYNOPSIS**43682 `#include <string.h>`43683 `int stricmp(const char *s1, const char *s2);`43684 **DESCRIPTION**

43685 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 43686 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43687 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

43688 The `stricmp()` function shall compare the string pointed to by `s1` to the string pointed to by `s2`.

43689 The sign of a non-zero return value shall be determined by the sign of the difference between the  
 43690 values of the first pair of bytes (both interpreted as type **unsigned char**) that differ in the strings  
 43691 being compared.

43692 **RETURN VALUE**

43693 Upon completion, `stricmp()` shall return an integer greater than, equal to, or less than 0, if the  
 43694 string pointed to by `s1` is greater than, equal to, or less than the string pointed to by `s2`,  
 43695 respectively.

43696 **ERRORS**

43697 No errors are defined.

43698 **EXAMPLES**43699 **Checking a Password Entry**

43700 The following example compares the information read from standard input to the value of the  
 43701 name of the user entry. If the `stricmp()` function returns 0 (indicating a match), a further check  
 43702 will be made to see if the user entered the proper old password. The `crypt()` function is used to  
 43703 encrypt the old password entered by the user, using the value of the encrypted password in the  
 43704 **passwd** structure as the salt. If this value matches the value of the encrypted **passwd** in the  
 43705 structure, the entered password `oldpasswd` is the correct user's password. Finally, the program  
 43706 encrypts the new password so that it can store the information in the **passwd** structure.

```

43707 #include <string.h>
43708 #include <unistd.h>
43709 #include <stdio.h>
43710 ...
43711 int valid_change;
43712 struct passwd *p;
43713 char user[100];
43714 char oldpasswd[100];
43715 char newpasswd[100];
43716 char savepasswd[100];
43717 ...
43718 if (stricmp(p->pw_name, user) == 0) {
43719 if (stricmp(p->pw_passwd, crypt(oldpasswd, p->pw_passwd)) == 0) {
43720 strcpy(savepasswd, crypt(newpasswd, user));
43721 p->pw_passwd = savepasswd;
43722 valid_change = 1;
43723 }
43724 else {

```

```
43725 fprintf(stderr, "Old password is not valid\n");
43726 }
43727 }
43728 ...
```

**43729 APPLICATION USAGE**

43730 None.

**43731 RATIONALE**

43732 None.

**43733 FUTURE DIRECTIONS**

43734 None.

**43735 SEE ALSO**

43736 *strncmp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>

**43737 CHANGE HISTORY**

43738 First released in Issue 1. Derived from Issue 1 of the SVID.

**43739 Issue 4**

43740 The DESCRIPTION is changed to make it clear that *strcmp()* compares bytes rather than  
43741 (possibly multi-byte) characters.

43742 The following change is incorporated for alignment with the ISO C standard:

- 43743 • The type of arguments *s1* and *s2* is changed from **char\*** to **const char\***.

**43744 Issue 6**

43745 Extensions beyond the ISO C standard are now marked.

43746 **NAME**

43747 strcoll — string comparison using collating information

43748 **SYNOPSIS**

43749 #include &lt;string.h&gt;

43750 int strcoll(const char \*s1, const char \*s2);

43751 **DESCRIPTION**

43752 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 43753 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43754 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

43755 The *strcoll()* function shall compare the string pointed to by *s1* to the string pointed to by *s2*,  
 43756 both interpreted as appropriate to the *LC\_COLLATE* category of the current locale.

43757 CX The *strcoll()* function shall not change the setting of *errno* if successful.

43758 Because no return value is reserved to indicate an error, an application wishing to check for error  
 43759 situations should set *errno* to 0, then call *strcoll()*, then check *errno*.

43760 **RETURN VALUE**

43761 Upon successful completion, *strcoll()* shall return an integer greater than, equal to, or less than 0,  
 43762 according to whether the string pointed to by *s1* is greater than, equal to, or less than the string  
 43763 CX pointed to by *s2* when both are interpreted as appropriate to the current locale. On error,  
 43764 *strcoll()* may set *errno*, but no return value is reserved to indicate an error.

43765 **ERRORS**43766 The *strcoll()* function may fail if:

43767 CX [EINVAL] The *s1* or *s2* arguments contain characters outside the domain of the collating  
 43768 sequence.

43769 **EXAMPLES**43770 **Comparing Nodes**

43771 The following example uses an application-defined function, *node\_compare()*, to compare two  
 43772 nodes based on an alphabetical ordering of the *string* field.

```
43773 #include <string.h>
43774 ...
43775 struct node { /* These are stored in the table. */
43776 char *string;
43777 int length;
43778 };
43779 ...
43780 int node_compare(const void *node1, const void *node2)
43781 {
43782 return strcoll(((const struct node *)node1)->string,
43783 ((const struct node *)node2)->string);
43784 }
43785 ...
```

43786 **APPLICATION USAGE**43787 The *strxfrm()* and *strcmp()* functions should be used for sorting large lists.

43788 **RATIONALE**

43789 None.

43790 **FUTURE DIRECTIONS**

43791 None.

43792 **SEE ALSO**43793 *strcmp()*, *strxfrm()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>43794 **CHANGE HISTORY**

43795 First released in Issue 3.

43796 **Issue 4**43797 A paragraph describing how the sign of the return value should be determined is removed from  
43798 the DESCRIPTION.

43799 The [EINVAL] error is marked as an extension.

43800 The following changes are incorporated for alignment with the ISO C standard:

- 43801 • The function is no longer marked as an extension.
- 43802 • The type of arguments *s1* and *s2* are changed from **char\*** to **const char\***.

43803 **Issue 5**43804 The DESCRIPTION is updated to indicate that *errno* does not be changed if the function is  
43805 successful.43806 **Issue 6**

43807 Extensions beyond the ISO C standard are now marked.

43808 The following new requirements on POSIX implementations derive from alignment with the  
43809 Single UNIX Specification:

- 43810 • The [EINVAL] optional error condition is added.

## 43811 NAME

43812 strcpy — copy a string

## 43813 SYNOPSIS

43814 #include &lt;string.h&gt;

43815 char \*strcpy(char \*restrict s1, const char \*restrict s2);

## 43816 DESCRIPTION

43817 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 43818 conflict between the requirements described here and the ISO C standard is unintentional. This  
 43819 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

43820 The *strcpy()* function shall copy the string pointed to by *s2* (including the terminating null byte)  
 43821 into the array pointed to by *s1*. If copying takes place between objects that overlap, the behavior  
 43822 is undefined.

## 43823 RETURN VALUE

43824 The *strcpy()* function shall return *s1*; no return value is reserved to indicate an error.

## 43825 ERRORS

43826 No errors are defined.

## 43827 EXAMPLES

43828 **Initializing a String**43829 The following example copies the string "-----" into the *permstring* variable.

```
43830 #include <string.h>
43831 ...
43832 static char permstring[11];
43833 ...
43834 strcpy(permstring, "-----");
43835 ...
```

43836 **Storing a Key and Data**

43837 The following example allocates space for a key using *malloc()* then uses *strcpy()* to place the  
 43838 key there. Then it allocates space for data using *malloc()*, and uses *strcpy()* to place data there.  
 43839 (The user-defined function *dbfree()* frees memory previously allocated to an array of type **struct**  
 43840 **element**.)

```
43841 #include <string.h>
43842 #include <stdlib.h>
43843 #include <stdio.h>
43844 ...
43845 /* Structure used to read data and store it. */
43846 struct element {
43847 char *key;
43848 char *data;
43849 };
43850 struct element *tbl, *curtbl;
43851 char *key, *data;
43852 int count;
43853 ...
43854 void dbfree(struct element *, int);
```

```

43855 ...
43856 if ((curtbl->key = malloc(strlen(key) + 1)) == NULL) {
43857 perror("malloc"); dbfree(tbl, count); return NULL;
43858 }
43859 strcpy(curtbl->key, key);

43860 if ((curtbl->data = malloc(strlen(data) + 1)) == NULL) {
43861 perror("malloc"); free(curtbl->key); dbfree(tbl, count); return NULL;
43862 }
43863 strcpy(curtbl->data, data);
43864 ...

```

#### 43865 APPLICATION USAGE

43866 Character movement is performed differently in different implementations. Thus, overlapping  
 43867 moves may yield surprises.

43868 This issue is aligned with the ISO C standard; this does not affect compatibility with XPG3  
 43869 applications. Reliable error detection by this function was never guaranteed.

#### 43870 RATIONALE

43871 None.

#### 43872 FUTURE DIRECTIONS

43873 None.

#### 43874 SEE ALSO

43875 *strncpy()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>

#### 43876 CHANGE HISTORY

43877 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 43878 Issue 4

43879 The DESCRIPTION is changed to make it clear that the function manipulates bytes rather than  
 43880 (possibly multi-byte) characters.

43881 The following change is incorporated for alignment with the ISO C standard:

- 43882 • The type of argument *s2* is changed from **char\*** to **const char\***.

#### 43883 Issue 6

43884 The *strcpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

43885 **NAME**

43886 strcspn — get length of a complementary substring

43887 **SYNOPSIS**

43888 #include &lt;string.h&gt;

43889 size\_t strcspn(const char \*s1, const char \*s2);

43890 **DESCRIPTION**

43891 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
43892 conflict between the requirements described here and the ISO C standard is unintentional. This  
43893 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

43894 The *strcspn()* function shall compute the length of the maximum initial segment of the string  
43895 pointed to by *s1* which consists entirely of bytes *not* from the string pointed to by *s2*.

43896 **RETURN VALUE**

43897 The *strcspn()* function shall return the length of the computed segment of the string pointed to  
43898 by *s1*; no return value is reserved to indicate an error.

43899 **ERRORS**

43900 No errors are defined.

43901 **EXAMPLES**

43902 None.

43903 **APPLICATION USAGE**

43904 None.

43905 **RATIONALE**

43906 None.

43907 **FUTURE DIRECTIONS**

43908 None.

43909 **SEE ALSO**43910 *strspn()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>43911 **CHANGE HISTORY**

43912 First released in Issue 1. Derived from Issue 1 of the SVID.

43913 **Issue 4**

43914 The DESCRIPTION is changed to make it clear that the function manipulates bytes rather than  
43915 (possibly multi-byte) characters.

43916 The following change is incorporated for alignment with the ISO C standard:

- 43917 • The type of arguments *s1* and *s2* is changed from **char\*** to **const char\***.

43918 **Issue 5**

43919 The RETURN VALUE section is updated to indicated that *strcspn()* returns the length of *s1*, and  
43920 not *s1* itself as was previously stated.

43921 **Issue 6**

43922 The Open Group corrigenda item U030/1 has been applied. The text of the RETURN VALUE  
43923 section is updated to indicate that the computed segment length is returned, not the *s1* length.

43924 **NAME**

43925            strdup — duplicate a string

43926 **SYNOPSIS**

43927 XSI        #include &lt;string.h&gt;

43928            char \*strdup(const char \*s1);

43929

43930 **DESCRIPTION**

43931            The *strdup()* function shall return a pointer to a new string, which is a duplicate of the string pointed to by *s1*. The returned pointer can be passed to *free()*. A null pointer is returned if the new string cannot be created.

43934 **RETURN VALUE**

43935            The *strdup()* function shall return a pointer to a new string on success. Otherwise, it shall return a null pointer and set *errno* to indicate the error.

43937 **ERRORS**43938            The *strdup()* function may fail if:

43939            [ENOMEM]       Storage space available is insufficient.

43940 **EXAMPLES**

43941            None.

43942 **APPLICATION USAGE**

43943            None.

43944 **RATIONALE**

43945            None.

43946 **FUTURE DIRECTIONS**

43947            None.

43948 **SEE ALSO**43949            *free()*, *malloc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>43950 **CHANGE HISTORY**

43951            First released in Issue 4, Version 2.

43952 **Issue 5**

43953            Moved from X/OPEN UNIX extension to BASE.

43954 **NAME**

43955            strerror, strerror\_r — get error message string

43956 **SYNOPSIS**

43957            #include &lt;string.h&gt;

43958            char \*strerror(int *errnum*);43959            int strerror\_r(int *errnum*, char \**strerrbuf*, size\_t *buflen*);43960 **DESCRIPTION**

43961            The *strerror()* function maps the error number in *errnum* to a locale-dependent error message string and returns a pointer to it. The string pointed to must not be modified by the application, but may be overwritten by a subsequent call to *strerror()* or *perror()*.

43962            The contents of the error message strings returned by *strerror()* should be determined by the setting of the *LC\_MESSAGES* category in the current locale.

43963            The implementation shall behave as if no function defined in this volume of IEEE Std. 1003.1-200x calls *strerror()*.

43964            The *strerror()* function shall not change the setting of *errno* if successful.

43965            Because no return value is reserved to indicate an error, an application wishing to check for error situations should set *errno* to 0, then call *strerror()*, then check *errno*.

43966            The *strerror()* function need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

43967            The *strerror\_r()* function maps the error number in *errnum* to a locale-dependent error message string and returns the string in the buffer pointed to by *strerrbuf*, with length *buflen*.

43975 **RETURN VALUE**

43976            Upon successful completion, *strerror()* shall return a pointer to the generated message string. On error *errno* may be set, but no return value is reserved to indicate an error.

43977            Upon successful completion, *strerror\_r()* returns 0. Otherwise, an error number is returned to indicate the error.

43980 **ERRORS**

43981            These functions may fail if:

43982            [EINVAL]            The value of *errnum* is not a valid error number.

43983            The *strerror\_r()* function may fail if:

43984            [ERANGE]            Insufficient storage was supplied via *strerrbuf* and *buflen* to contain the generated message string.

43986 **EXAMPLES**

43987            None.

43988 **APPLICATION USAGE**

43989            None.

43990 **RATIONALE**

43991            None.

43992 **FUTURE DIRECTIONS**

43993            None.

43994 **SEE ALSO**

43995 *perror()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>

43996 **CHANGE HISTORY**

43997 First released in Issue 3.

43998 **Issue 4**

43999 The DESCRIPTION is changed as follows:

- 44000 • The term “language-dependent” is replaced by “locale-dependent”.
- 44001 • A statement about the use of the *LC\_MESSAGES* category for determining the language of
- 44002 error messages is added and marked as an extension.

44003 The fact that *strerror()* can return a null pointer on failure and set *errno* is marked as an

44004 extension.

44005 The [EINVAL] error is marked as an extension.

44006 The FUTURE DIRECTIONS section is removed.

44007 The following change is incorporated for alignment with the ISO C standard:

- 44008 • The function is no longer marked as an extension.

44009 **Issue 5**

44010 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

44011 A note indicating that this function need not be reentrant is added to the DESCRIPTION.

44012 **Issue 6**

44013 Extensions beyond the ISO C standard are now marked.

44014 The following new requirements on POSIX implementations derive from alignment with the

44015 Single UNIX Specification:

- 44016 • In the RETURN VALUE section, the fact that *errno* may be set is added.
- 44017 • The [EINVAL] optional error condition is added.

44018 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

44019 The *strerror\_r()* function is added in response to IEEE PASC Interpretation 1003.1c #39.

## 44020 NAME

44021 strfmon — convert monetary value to a string

## 44022 SYNOPSIS

44023 XSI 

```
#include <monetary.h>
```

44024 

```
ssize_t strfmon(char *restrict s, size_t maxsize,
44025 const char *restrict format, ...);
```

44026

## 44027 DESCRIPTION

44028 The *strfmon()* function shall place characters into the array pointed to by *s* as controlled by the  
44029 string pointed to by *format*. No more than *maxsize* bytes are placed into the array.44030 The format is a character string that contains two types of objects: *plain characters*, which are  
44031 simply copied to the output stream, and *conversion specifications*, each of which results in the  
44032 fetching of zero or more arguments which are converted and formatted. The results are  
44033 undefined if there are insufficient arguments for the format. If the format is exhausted while  
44034 arguments remain, the excess arguments are simply ignored.

44035 The application shall ensure that a conversion specification consists of the following sequence:

- 44036 • A '%' character
- 44037 • Optional flags
- 44038 • Optional field width
- 44039 • Optional left precision
- 44040 • Optional right precision
- 44041 • A required conversion character that determines the conversion to be performed

44042 **Flags**

44043 One or more of the following optional flags can be specified to control the conversion:

- 44044 =*f* An '=' followed by a single character *f* which is used as the numeric fill character. The  
44045 application shall ensure that the fill character is representable in a single byte in order  
44046 to work with precision and width counts. The default numeric fill character is the  
44047 <space> character. This flag does not affect field width filling which always uses the  
44048 <space> character. This flag is ignored unless a left precision (see below) is specified.
- 44049 ^ Do not format the currency amount with grouping characters. The default is to insert  
44050 the grouping characters if defined for the current locale.
- 44051 + or ( Specify the style of representing positive and negative currency amounts. Only one of  
44052 '+' or '(' may be specified. If '+' is specified, the locale's equivalent of '+' and '-'  
44053 are used (for example, in the U.S., the empty string if positive and '-' if negative). If  
44054 '(' is specified, negative amounts are enclosed within parentheses. If neither flag is  
44055 specified, the '+' style is used.
- 44056 ! Suppress the currency symbol from the output conversion.
- 44057 - Specify the alignment. If this flag is present all fields are left-justified (padded to the  
44058 right) rather than right-justified.

44059 **Field Width**

44060 *w* A decimal digit string *w* specifying a minimum field width in bytes in which the result  
 44061 of the conversion is right-justified (or left-justified if the flag '-' is specified). The  
 44062 default is 0.

44063 **Left Precision**

44064 *#n* A '#' followed by a decimal digit string *n* specifying a maximum number of digits  
 44065 expected to be formatted to the left of the radix character. This option can be used to  
 44066 keep the formatted output from multiple calls to the *strfmon()* function aligned in the  
 44067 same columns. It can also be used to fill unused positions with a special character as in  
 44068 "\$\*\*\*123.45". This option causes an amount to be formatted as if it has the number  
 44069 of digits specified by *n*. If more than *n* digit positions are required, this conversion  
 44070 specification is ignored. Digit positions in excess of those actually required are filled  
 44071 with the numeric fill character (see the *=f* flag above).

44072 If grouping has not been suppressed with the '^' flag, and it is defined for the current  
 44073 locale, grouping separators are inserted before the fill characters (if any) are added.  
 44074 Grouping separators are not applied to fill characters even if the fill character is a digit.

44075 To ensure alignment, any characters appearing before or after the number in the  
 44076 formatted output such as currency or sign symbols are padded as necessary with  
 44077 <space> characters to make their positive and negative formats an equal length.

44078 **Right Precision**

44079 *.p* A period followed by a decimal digit string *p* specifying the number of digits after the  
 44080 radix character. If the value of the right precision *p* is 0, no radix character appears. If a  
 44081 right precision is not included, a default specified by the current locale is used. The  
 44082 amount being formatted is rounded to the specified number of digits prior to  
 44083 formatting.

44084 **Conversion Characters**

44085 The conversion characters and their meanings are:

44086 *i* The **double** argument is formatted according to the locale's international currency  
 44087 format (for example, in the U.S.: USD 1,234.56).

44088 *n* The **double** argument is formatted according to the locale's national currency format  
 44089 (for example, in the U.S.: \$1,234.56).

44090 *%* Convert to a '%'; no argument is converted. The entire conversion specification shall  
 44091 be "%%".

44092 **Locale Information**

44093 The *LC\_MONETARY* category of the program's locale affects the behavior of this function  
 44094 including the monetary radix character (which may be different from the numeric radix  
 44095 character affected by the *LC\_NUMERIC* category), the grouping separator, the currency  
 44096 symbols, and formats. The international currency symbol should be conformant with the  
 44097 ISO 4217:1995 standard.

44098 If the value of *maxsize* is greater than {SSIZE\_MAX}, the result is implementation-defined.

## 44099 RETURN VALUE

44100 If the total number of resulting bytes including the terminating null byte is not more than  
 44101 *maxsize*, *strfmon()* shall return the number of bytes placed into the array pointed to by *s*, not  
 44102 including the terminating null byte. Otherwise, -1 shall be returned, the contents of the array are  
 44103 indeterminate, and *errno* shall be set to indicate the error.

## 44104 ERRORS

44105 The *strfmon()* function shall fail if:

44106 [E2BIG] Conversion stopped due to lack of space in the buffer.

## 44107 EXAMPLES

44108 Given a locale for the U.S. and the values 123.45, -123.45, and 3456.781:

| Conversion Specification | Output                                      | Comments                                                                   |
|--------------------------|---------------------------------------------|----------------------------------------------------------------------------|
| %n                       | \$123.45<br>-\$123.45<br>\$3,456.78         | Default formatting.                                                        |
| %11n                     | \$123.45<br>-\$123.45<br>\$3,456.78         | Right align within an 11 character field.                                  |
| ##5n                     | \$ 123.45<br>-\$ 123.45<br>\$ 3,456.78      | Aligned columns for values up to 99,999.                                   |
| %=*#5n                   | \$***123.45<br>-\$***123.45<br>\$*3,456.78  | Specify a fill character.                                                  |
| %=#5n                    | \$000123.45<br>-\$000123.45<br>\$03,456.78  | Fill characters do not use grouping even if the fill character is a digit. |
| %^#5n                    | \$ 123.45<br>-\$ 123.45<br>\$ 3456.78       | Disable the grouping separator.                                            |
| %^#5.0n                  | \$ 123<br>-\$ 123<br>\$ 3457                | Round off to whole units.                                                  |
| %^#5.4n                  | \$ 123.4500<br>-\$ 123.4500<br>\$ 3456.7810 | Increase the precision.                                                    |
| %(#5n                    | 123.45<br>(\$ 123.45)<br>\$ 3,456.78        | Use an alternative pos/neg style.                                          |
| %(!#5n                   | 123.45<br>( 123.45)<br>3,456.78             | Disable the currency symbol.                                               |

## 44141 APPLICATION USAGE

44142 None.

44143 **RATIONALE**

44144 None.

44145 **FUTURE DIRECTIONS**44146 Lowercase conversion characters are reserved for future standards use and uppercase for  
44147 implementation-defined use.44148 **SEE ALSO**44149 *localeconv()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**monetary.h**>44150 **CHANGE HISTORY**

44151 First released in Issue 4.

44152 **Issue 5**

44153 Moved from ENHANCED I18N to BASE.

44154 The [ENOSYS] error is removed.

44155 A sentence is added to the DESCRIPTION warning about values of *maxsize* that are greater than  
44156 {SSIZE\_MAX}.44157 **Issue 6**

44158 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

44159 The **restrict** keyword is added to the *strfmon()* prototype for alignment with the  
44160 ISO/IEC 9899:1999 standard.

## 44161 NAME

44162 strftime — convert date and time to a string

## 44163 SYNOPSIS

44164 #include &lt;time.h&gt;

```
44165 size_t strftime(char *restrict s, size_t maxsize,
44166 const char *restrict format, const struct tm *restrict timptr);
```

## 44167 DESCRIPTION

44168 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 44169 conflict between the requirements described here and the ISO C standard is unintentional. This  
 44170 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44171 The *strftime()* function shall place bytes into the array pointed to by *s* as controlled by the string  
 44172 pointed to by *format*. The *format* string consists of zero or more conversion specifications and  
 44173 ordinary characters. A conversion specification consists of a '%' character, possibly followed by  
 44174 an *E* or *O* modifier, and a terminating conversion character that determines the conversion  
 44175 specification's behavior. All ordinary characters (including the terminating null byte) are copied  
 44176 unchanged into the array. If copying takes place between objects that overlap, the behavior is  
 44177 undefined. No more than *maxsize* bytes are placed into the array. Each conversion specification  
 44178 is replaced by appropriate characters as described in the following list. The appropriate  
 44179 characters are determined by the program's locale and by the values contained in the structure  
 44180 pointed to by *timptr*.

44181 cx Local timezone information is used as though *strftime()* called *tzset()*.

|       |    |                                                                                         |
|-------|----|-----------------------------------------------------------------------------------------|
| 44182 | %a | Replaced by the locale's abbreviated weekday name.                                      |
| 44183 | %A | Replaced by the locale's full weekday name.                                             |
| 44184 | %b | Replaced by the locale's abbreviated month name.                                        |
| 44185 | %B | Replaced by the locale's full month name.                                               |
| 44186 | %c | Replaced by the locale's appropriate date and time representation.                      |
| 44187 | %C | Replaced by the century number (the year divided by 100 and truncated to an integer)    |
| 44188 |    | as a decimal number [00-99].                                                            |
| 44189 | %d | Replaced by the day of the month as a decimal number [01,31].                           |
| 44190 | %D | Same as %m/%d/%y.                                                                       |
| 44191 | %e | Replaced by the day of the month as a decimal number [1,31]; a single digit is preceded |
| 44192 |    | by a space.                                                                             |
| 44193 | %F | Equivalent to %Y-%m-%d (the ISO 8601: 1988 standard date format).                       |
| 44194 | %g | Replaced by the last 2 digits of the week-based year (see below) as a decimal number    |
| 44195 |    | (00-99).                                                                                |
| 44196 | %G | Replaced by the week-based year (see below) as a decimal number (for example, 1977).    |
| 44197 | %h | Same as %b.                                                                             |
| 44198 | %H | Replaced by the hour (24-hour clock) as a decimal number [00,23].                       |
| 44199 | %I | Replaced by the hour (12-hour clock) as a decimal number [01,12].                       |
| 44200 | %j | Replaced by the day of the year as a decimal number [001,366].                          |

|       |                 |                                                                                                   |
|-------|-----------------|---------------------------------------------------------------------------------------------------|
| 44201 | <code>%m</code> | Replaced by the month as a decimal number [01,12].                                                |
| 44202 | <code>%M</code> | Replaced by the minute as a decimal number [00,59].                                               |
| 44203 | <code>%n</code> | Replaced by a <newline> character.                                                                |
| 44204 | <code>%p</code> | Replaced by the locale's equivalent of either a.m. or p.m.                                        |
| 44205 | <code>%r</code> | Replaced by the time in a.m. and p.m. notation; in the POSIX locale this is equivalent to         |
| 44206 |                 | <code>%I:%M:%S %p</code> .                                                                        |
| 44207 | <code>%R</code> | Replaced by the time in 24 hour notation ( <code>%H:%M</code> ).                                  |
| 44208 | <code>%S</code> | Replaced by the second as a decimal number [00,61].                                               |
| 44209 | <code>%t</code> | Replaced by a <tab> character.                                                                    |
| 44210 | <code>%T</code> | Replaced by the time ( <code>%H:%M:%S</code> ).                                                   |
| 44211 | <code>%u</code> | Replaced by the weekday as a decimal number [1,7], with 1 representing Monday.                    |
| 44212 | <code>%U</code> | Replaced by the week number of the year (Sunday as the first day of the week) as a                |
| 44213 |                 | decimal number [00,53].                                                                           |
| 44214 | <code>%V</code> | Replaced by the week number of the year (Monday as the first day of the week) as a                |
| 44215 |                 | decimal number [01,53]. If the week containing 1 January has four or more days in the             |
| 44216 |                 | new year, then it is considered week 1. Otherwise, it is the last week of the previous            |
| 44217 |                 | year, and the next week is week 1.                                                                |
| 44218 | <code>%w</code> | Replaced by the weekday as a decimal number [0,6], with 0 representing Sunday.                    |
| 44219 | <code>%W</code> | Replaced by the week number of the year (Monday as the first day of the week) as a                |
| 44220 |                 | decimal number [00,53]. All days in a new year preceding the first Monday are                     |
| 44221 |                 | considered to be in week 0.                                                                       |
| 44222 | <code>%x</code> | Replaced by the locale's appropriate date representation.                                         |
| 44223 | <code>%X</code> | Replaced by the locale's appropriate time representation.                                         |
| 44224 | <code>%y</code> | Replaced by the year without century as a decimal number [00,99].                                 |
| 44225 | <code>%Y</code> | Replaced by the year with century as a decimal number.                                            |
| 44226 | <code>%z</code> | Replaced by the offset from UTC in the ISO 8601:1988 standard format <code>-0430</code> (meaning  |
| 44227 |                 | 4 hours 30 minutes behind UTC, west of Greenwich), or by no characters if no timezone             |
| 44228 |                 | is determinable.                                                                                  |
| 44229 | <code>%Z</code> | Replaced by the timezone name or abbreviation, or by no bytes if no timezone                      |
| 44230 |                 | information exists.                                                                               |
| 44231 | <code>%%</code> | Replaced by <code>'%'</code> .                                                                    |
| 44232 |                 | If a conversion specification does not correspond to any of the above, the behavior is undefined. |

#### 44233 **Modified Conversion Specifiers**

|       |                  |                                                                                                                |
|-------|------------------|----------------------------------------------------------------------------------------------------------------|
| 44234 | <code>CX</code>  | Some conversion specifiers can be modified by the <i>E</i> or <i>O</i> modifier characters to indicate that an |
| 44235 |                  | alternative format or specification should be used rather than the one normally used by the                    |
| 44236 |                  | unmodified conversion specifier. If the alternative format or specification does not exist for the             |
| 44237 |                  | current locale, (see ERA in the Base Definitions volume of IEEE Std. 1003.1-200x, Section 7.3.5,               |
| 44238 |                  | LC_TIME) the behavior shall be as if the unmodified conversion specification were used.                        |
| 44239 | <code>%Ec</code> | Replaced by the locale's alternative appropriate date and time representation.                                 |

|       |            |                                                                                                                                                                                                    |
|-------|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 44240 | <b>%EC</b> | Replaced by the name of the base year (period) in the locale's alternative representation.                                                                                                         |
| 44241 |            |                                                                                                                                                                                                    |
| 44242 | <b>%Ex</b> | Replaced by the locale's alternative date representation.                                                                                                                                          |
| 44243 | <b>%EX</b> | Replaced by the locale's alternative time representation.                                                                                                                                          |
| 44244 | <b>%Ey</b> | Replaced by the offset from <b>%EC</b> (year only) in the locale's alternative representation.                                                                                                     |
| 44245 | <b>%EY</b> | Replaced by the full alternative year representation.                                                                                                                                              |
| 44246 | <b>%Od</b> | Replaced by the day of the month, using the locale's alternative numeric symbols, filled as needed with leading zeros if there is any alternative symbol for zero; otherwise, with leading spaces. |
| 44247 |            |                                                                                                                                                                                                    |
| 44248 |            |                                                                                                                                                                                                    |
| 44249 | <b>%Oe</b> | Replaced by the day of the month, using the locale's alternative numeric symbols, filled as needed with leading spaces.                                                                            |
| 44250 |            |                                                                                                                                                                                                    |
| 44251 | <b>%OH</b> | Replaced by the hour (24-hour clock) using the locale's alternative numeric symbols.                                                                                                               |
| 44252 | <b>%OI</b> | Replaced by the hour (12-hour clock) using the locale's alternative numeric symbols.                                                                                                               |
| 44253 | <b>%Om</b> | Replaced by the month using the locale's alternative numeric symbols.                                                                                                                              |
| 44254 | <b>%OM</b> | Replaced by the minutes using the locale's alternative numeric symbols.                                                                                                                            |
| 44255 | <b>%OS</b> | Replaced by the seconds using the locale's alternative numeric symbols.                                                                                                                            |
| 44256 | <b>%Ou</b> | Replaced by the weekday as a number in the locale's alternative representation (Monday=1).                                                                                                         |
| 44257 |            |                                                                                                                                                                                                    |
| 44258 | <b>%OU</b> | Replaced by the week number of the year (Sunday as the first day of the week, rules corresponding to <b>%U</b> ) using the locale's alternative numeric symbols.                                   |
| 44259 |            |                                                                                                                                                                                                    |
| 44260 | <b>%OV</b> | Replaced by the week number of the year (Monday as the first day of the week, rules corresponding to <b>%V</b> ) using the locale's alternative numeric symbols.                                   |
| 44261 |            |                                                                                                                                                                                                    |
| 44262 | <b>%Ow</b> | Replaced by the number of the weekday (Sunday=0) using the locale's alternative numeric symbols.                                                                                                   |
| 44263 |            |                                                                                                                                                                                                    |
| 44264 | <b>%OW</b> | Replaced by the week number of the year (Monday as the first day of the week) using the locale's alternative numeric symbols.                                                                      |
| 44265 |            |                                                                                                                                                                                                    |
| 44266 | <b>%Oy</b> | Replaced by the year (offset from <b>%C</b> ) using the locale's alternative numeric symbols.                                                                                                      |
| 44267 |            |                                                                                                                                                                                                    |
| 44268 |            |                                                                                                                                                                                                    |
| 44269 |            |                                                                                                                                                                                                    |
| 44270 |            |                                                                                                                                                                                                    |
| 44271 |            |                                                                                                                                                                                                    |
| 44272 |            |                                                                                                                                                                                                    |
| 44273 |            |                                                                                                                                                                                                    |
| 44274 |            |                                                                                                                                                                                                    |
| 44275 |            |                                                                                                                                                                                                    |
| 44276 |            |                                                                                                                                                                                                    |
| 44277 |            |                                                                                                                                                                                                    |
| 44278 |            |                                                                                                                                                                                                    |
| 44279 |            |                                                                                                                                                                                                    |
| 44280 |            |                                                                                                                                                                                                    |

44281 **ERRORS**

44282 No errors are defined.

44283 **EXAMPLES**44284 **Getting a Localized Date String**

44285 The following example first sets the locale to the user's default. The locale information will be  
 44286 used in the `nl_langinfo()` and `strptime()` functions. The `nl_langinfo()` function returns the localized  
 44287 date string which specifies how the date is laid out. The `strptime()` function takes this information  
 44288 and, using the `tm` structure for values, places the date and time information into `datestring`.

```
44289 #include <time.h>
44290 #include <locale.h>
44291 #include <langinfo.h>
44292 ...
44293 struct tm *tm;
44294 char datestring[256];
44295 ...
44296 setlocale (LC_ALL, "");
44297 ...
44298 strptime (datestring, sizeof(datestring), nl_langinfo (D_T_FMT), tm);
44299 ...
```

44300 **APPLICATION USAGE**

44301 The range of values for `%S` is [00,61] rather than [00,59] to allow for the occasional leap second  
 44302 and even more infrequent double leap second.

44303 Some of the conversion specifications marked EX are duplicates of others. They are included for  
 44304 compatibility with `nl_cxtime()` and `nl_ascxtime()`, which were published in Issue 2.

44305 Applications should use `%Y` (4-digit years) in preference to `%y` (2-digit years).

44306 In the C locale, the `E` and `O` modifiers are ignored and the replacement strings for the following  
 44307 specifiers are:

|       |                 |                                                 |
|-------|-----------------|-------------------------------------------------|
| 44308 | <code>%a</code> | The first three characters of <code>%A</code> . |
| 44309 | <code>%A</code> | One of Sunday, Monday, ..., Saturday.           |
| 44310 | <code>%b</code> | The first three characters of <code>%B</code> . |
| 44311 | <code>%B</code> | One of January, February, ..., December.        |
| 44312 | <code>%c</code> | Equivalent to <code>%a %b %e %T %Y</code> .     |
| 44313 | <code>%p</code> | One of AM or PM.                                |
| 44314 | <code>%r</code> | Equivalent to <code>%I:%M:%S %p</code> .        |
| 44315 | <code>%x</code> | Equivalent to <code>%m/%d/%y</code> .           |
| 44316 | <code>%X</code> | Equivalent to <code>%T</code> .                 |
| 44317 | <code>%Z</code> | Implementation-defined.                         |

44318 **RATIONALE**

44319 None.

44320 **FUTURE DIRECTIONS**

44321 None.

44322 **SEE ALSO**44323 *asctime()*, *clock()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strptime()*, *time()*, *utime()*,  
44324 the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>44325 **CHANGE HISTORY**

44326 First released in Issue 3.

44327 **Issue 4**

44328 The DESCRIPTION is expanded to describe modified conversion specifiers.

44329 %C, %e, %R, %u, and %V are added to the list of valid conversion specifications.

44330 The DESCRIPTION and RETURN VALUE sections are changed to make it clear when the  
44331 function uses byte values rather than (possibly multi-byte) character values.

44332 The following changes are incorporated for alignment with the ISO C standard:

- 44333
- The type of argument *format* is changed from **char\*** to **const char\***, and the type of argument  
44334 *tm\_ptr* is changed from **struct tm\*** to **const struct tm\***.
  - In the description of the %Z conversion specification, the words “or abbreviation” are added  
44335 to indicate that *strptime()* does not necessarily return a full timezone name.  
44336

44337 **Issue 5**44338 The description of %OV is changed to be consistent with %V and defines Monday as the first  
44339 day of the week.

44340 The description of %Oy is clarified.

44341 **Issue 6**

44342 Extensions beyond the ISO C standard are now marked.

44343 The Open Group corrigenda item U033/8 has been applied. The %V conversion specifier is  
44344 changed from “Otherwise, it is week 53 of the previous year, and the next week is week 1” to  
44345 “Otherwise, it is the last week of the previous year, and the next week is week 1”.44346 The following new requirements on POSIX implementations derive from alignment with the  
44347 Single UNIX Specification:

- 44348
- The %C, %D, %e, %h, %n, %r, %R, %t, and %T conversion specifiers are added.
  - The modified conversion specifiers are added for consistency with the ISO POSIX-2 standard  
44349 *date* utility.  
44350

44351 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 44352
- The *strptime()* prototype is updated.
  - The DESCRIPTION is extensively revised.  
44353

44354 **NAME**

44355            strlen — get string length

44356 **SYNOPSIS**

44357            #include &lt;string.h&gt;

44358            size\_t strlen(const char \*s);

44359 **DESCRIPTION**

44360 cx        The functionality described on this reference page is aligned with the ISO C standard. Any  
44361        conflict between the requirements described here and the ISO C standard is unintentional. This  
44362        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44363        The *strlen()* function shall compute the number of bytes in the string to which *s* points, not  
44364        including the terminating null byte.

44365 **RETURN VALUE**

44366        The *strlen()* function shall return the length of *s*; no return value shall be reserved to indicate an  
44367        error.

44368 **ERRORS**

44369        No errors are defined.

44370 **EXAMPLES**44371            **Getting String Lengths**

44372        The following example sets the maximum length of *key* and *data* by using *strlen()* to get the  
44373        lengths of those strings.

```
44374 #include <string.h>
44375 ...
44376 struct element {
44377 char *key;
44378 char *data;
44379 };
44380 ...
44381 char *key, *data;
44382 int len;
44383 *keylength = *datalength = 0;
44384 ...
44385 if ((len = strlen(key)) > *keylength)
44386 *keylength = len;
44387 if ((len = strlen(data)) > *datalength)
44388 *datalength = len;
44389 ...
```

44390 **APPLICATION USAGE**

44391        None.

44392 **RATIONALE**

44393        None.

44394 **FUTURE DIRECTIONS**

44395        None.

44396 **SEE ALSO**

44397 The Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>

44398 **CHANGE HISTORY**

44399 First released in Issue 1. Derived from Issue 1 of the SVID.

44400 **Issue 4**

44401 The DESCRIPTION is changed to make it clear that the function works in units of bytes rather than (possibly multi-byte) characters.

44403 The following change is incorporated for alignment with the ISO C standard:

- 44404 • The type of argument *s* is changed from **char\*** to **const char\***.

44405 **Issue 5**

44406 The RETURN VALUE section is updated to indicate that *strlen()* returns the length of *s*, and not *s* itself as was previously stated.

44407

44408 **NAME**

44409           strncasecmp — case-insensitive string comparison

44410 **SYNOPSIS**

44411 xSI       #include &lt;strings.h&gt;

44412           int strncasecmp(const char \*s1, const char \*s2, size\_t n);

44413

44414 **DESCRIPTION**44415           Refer to *strcasecmp()*.

44416 **NAME**

44417           strncat — concatenate a string with part of another

44418 **SYNOPSIS**

44419           #include <string.h>

44420           char \*strncat(char \*restrict *s1*, const char \*restrict *s2*, size\_t *n*);

44421 **DESCRIPTION**

44422 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
44423 conflict between the requirements described here and the ISO C standard is unintentional. This  
44424 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44425       The *strncat()* function shall append not more than *n* bytes (a null byte and bytes that follow it  
44426 are not appended) from the array pointed to by *s2* to the end of the string pointed to by *s1*. The  
44427 initial byte of *s2* overwrites the null byte at the end of *s1*. A terminating null byte is always  
44428 appended to the result. If copying takes place between objects that overlap, the behavior is  
44429 undefined.

44430 **RETURN VALUE**

44431       The *strncat()* function shall return *s1*; no return value shall be reserved to indicate an error.

44432 **ERRORS**

44433       No errors are defined.

44434 **EXAMPLES**

44435       None.

44436 **APPLICATION USAGE**

44437       None.

44438 **RATIONALE**

44439       None.

44440 **FUTURE DIRECTIONS**

44441       None.

44442 **SEE ALSO**

44443       *strcat()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>

44444 **CHANGE HISTORY**

44445       First released in Issue 1. Derived from Issue 1 of the SVID.

44446 **Issue 4**

44447       The **DESCRIPTION** is changed to make it clear that the function manipulates bytes rather than  
44448 (possibly multi-byte) characters.

44449       The following change is incorporated for alignment with the ISO C standard:

- 44450       • The type of argument *s2* is changed from **char\*** to **const char\***.

44451 **Issue 6**

44452       The *strncat()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

44453 **NAME**

44454 strncmp — compare part of two strings

44455 **SYNOPSIS**

44456 #include &lt;string.h&gt;

44457 int strncmp(const char \*s1, const char \*s2, size\_t n);

44458 **DESCRIPTION**

44459 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
44460 conflict between the requirements described here and the ISO C standard is unintentional. This  
44461 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44462 The *strncmp()* function shall compare not more than *n* bytes (bytes that follow a null byte are not  
44463 compared) from the array pointed to by *s1* to the array pointed to by *s2*.

44464 The sign of a non-zero return value is determined by the sign of the difference between the  
44465 values of the first pair of bytes (both interpreted as type **unsigned char**) that differ in the strings  
44466 being compared.

44467 **RETURN VALUE**

44468 Upon successful completion, *strncmp()* shall return an integer greater than, equal to, or less than  
44469 0, if the possibly null-terminated array pointed to by *s1* is greater than, equal to, or less than the  
44470 possibly null-terminated array pointed to by *s2* respectively.

44471 **ERRORS**

44472 No errors are defined.

44473 **EXAMPLES**

44474 None.

44475 **APPLICATION USAGE**

44476 None.

44477 **RATIONALE**

44478 None.

44479 **FUTURE DIRECTIONS**

44480 None.

44481 **SEE ALSO**44482 *strcmp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>44483 **CHANGE HISTORY**

44484 First released in Issue 1. Derived from Issue 1 of the SVID.

44485 **Issue 4**

44486 The DESCRIPTION is changed to make it clear that the function manipulates bytes rather than  
44487 (possibly multi-byte) characters.

44488 The following change is incorporated for alignment with the ISO C standard:

- 44489 • The type of arguments *s1* and *s2* are changed from **char\*** to **const char\***.

44490 **Issue 6**

44491 Extensions beyond the ISO C standard are now marked.

44492 **NAME**

44493       strncpy — copy part of a string

44494 **SYNOPSIS**

44495       #include &lt;string.h&gt;

44496       char \*strncpy(char \*restrict *s1*, const char \*restrict *s2*, size\_t *n*);44497 **DESCRIPTION**

44498 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
44499 conflict between the requirements described here and the ISO C standard is unintentional. This  
44500 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44501       The *strncpy()* function shall copy not more than *n* bytes (bytes that follow a null byte are not  
44502 copied) from the array pointed to by *s2* to the array pointed to by *s1*. If copying takes place  
44503 between objects that overlap, the behavior is undefined.

44504       If the array pointed to by *s2* is a string that is shorter than *n* bytes, null bytes shall be appended  
44505 to the copy in the array pointed to by *s1*, until *n* bytes in all are written.

44506 **RETURN VALUE**44507       The *strncpy()* function shall return *s1*; no return value is reserved to indicate an error.44508 **ERRORS**

44509       No errors are defined.

44510 **EXAMPLES**

44511       None.

44512 **APPLICATION USAGE**

44513       Character movement is performed differently in different implementations. Thus, overlapping  
44514 moves may yield surprises.

44515       If there is no null byte in the first *n* bytes of the array pointed to by *s2*, the result is not null-  
44516 terminated.

44517 **RATIONALE**

44518       None.

44519 **FUTURE DIRECTIONS**

44520       None.

44521 **SEE ALSO**44522       *strcpy()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>44523 **CHANGE HISTORY**

44524       First released in Issue 1. Derived from Issue 1 of the SVID.

44525 **Issue 4**

44526       The **DESCRIPTION** is changed to make it clear that the function manipulates bytes rather than  
44527 (possibly multi-byte) characters.

44528       The following change is incorporated for alignment with the ISO C standard:

- 44529       • The type of argument *s2* is changed from **char\*** to **const char\***.

44530 **Issue 6**44531       The *strncpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

44532 **NAME**

44533 strpbrk — scan string for byte

44534 **SYNOPSIS**

44535 #include &lt;string.h&gt;

44536 char \*strpbrk(const char \*s1, const char \*s2);

44537 **DESCRIPTION**

44538 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
44539 conflict between the requirements described here and the ISO C standard is unintentional. This  
44540 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44541 The *strpbrk()* function shall locate the first occurrence in the string pointed to by *s1* of any byte  
44542 from the string pointed to by *s2*.

44543 **RETURN VALUE**

44544 Upon successful completion, *strpbrk()* shall return a pointer to the byte or a null pointer if no  
44545 byte from *s2* occurs in *s1*.

44546 **ERRORS**

44547 No errors are defined.

44548 **EXAMPLES**

44549 None.

44550 **APPLICATION USAGE**

44551 None.

44552 **RATIONALE**

44553 None.

44554 **FUTURE DIRECTIONS**

44555 None.

44556 **SEE ALSO**44557 *strchr()*, *strchr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>44558 **CHANGE HISTORY**

44559 First released in Issue 1. Derived from Issue 1 of the SVID.

44560 **Issue 4**

44561 The DESCRIPTION and RETURN VALUE sections are changed to make it clear that the function  
44562 works in units of bytes rather than (possibly multi-byte) characters.

44563 The following change is incorporated for alignment with the ISO C standard:

- 44564 • The type of arguments *s1* and *s2* is changed from **char\*** to **const char\***.

## 44565 NAME

44566 strptime — date and time conversion

## 44567 SYNOPSIS

44568 XSI 

```
#include <time.h>
```

```
44569 char *strptime(const char *restrict buf, const char *restrict format,
44570 struct tm *restrict tm);
44571
```

## 44572 DESCRIPTION

44573 The *strptime()* function shall convert the character string pointed to by *buf* to values which are  
 44574 stored in the **tm** structure pointed to by *tm*, using the format specified by *format*.

44575 The *format* is composed of zero or more directives. Each directive is composed of one of the  
 44576 following: one or more white-space characters (as specified by *isspace()*); an ordinary character  
 44577 (neither '%' nor a white-space character); or a conversion specification. Each conversion  
 44578 specification is composed of a '%' character followed by a conversion character which specifies  
 44579 the replacement required. The application shall ensure that there is white-space or other non-  
 44580 alphanumeric characters between any two conversion specifications. The following conversion  
 44581 specifications are supported:

|       |    |                                                                                                                                                                                               |
|-------|----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 44582 | %a | The day of the week, using the locale's weekday names; either the abbreviated or full name may be specified.                                                                                  |
| 44583 |    |                                                                                                                                                                                               |
| 44584 | %A | The same as %a.                                                                                                                                                                               |
| 44585 | %b | The month, using the locale's month names; either the abbreviated or full name may be specified.                                                                                              |
| 44586 |    |                                                                                                                                                                                               |
| 44587 | %B | The same as %b.                                                                                                                                                                               |
| 44588 | %c | Replaced by the locale's appropriate date and time representation.                                                                                                                            |
| 44589 | %C | The century number [0,99]; leading zeros are permitted but not required.                                                                                                                      |
| 44590 | %d | The day of the month [1,31]; leading zeros are permitted but not required.                                                                                                                    |
| 44591 | %D | The date as %m/%d/%y.                                                                                                                                                                         |
| 44592 | %e | The same as %d.                                                                                                                                                                               |
| 44593 | %h | The same as %b.                                                                                                                                                                               |
| 44594 | %H | The hour (24-hour clock) [0,23]; leading zeros are permitted but not required.                                                                                                                |
| 44595 | %I | The hour (12-hour clock) [1,12]; leading zeros are permitted but not required.                                                                                                                |
| 44596 | %j | The day number of the year [1,366]; leading zeros are permitted but not required.                                                                                                             |
| 44597 | %m | The month number [1,12]; leading zeros are permitted but not required.                                                                                                                        |
| 44598 | %M | The minute [0-59]; leading zeros are permitted but not required.                                                                                                                              |
| 44599 | %n | Any white space.                                                                                                                                                                              |
| 44600 | %p | The locale's equivalent of a.m or p.m.                                                                                                                                                        |
| 44601 | %r | 12-hour clock time using the AM/PM notation if <b>t_fmt_ampm</b> is not an empty string in the LC_TIME portion of the current locale; in the POSIX locale, this is equivalent to %I:%M:%S %p. |
| 44602 |    |                                                                                                                                                                                               |
| 44603 |    |                                                                                                                                                                                               |
| 44604 | %R | The time as %H:%M.                                                                                                                                                                            |

|       |                  |                                                                                                                                                                                                                                                                                                                                                                         |
|-------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 44605 | <code>%S</code>  | The seconds [0,61]; leading zeros are permitted but not required.                                                                                                                                                                                                                                                                                                       |
| 44606 | <code>%t</code>  | Any white space.                                                                                                                                                                                                                                                                                                                                                        |
| 44607 | <code>%T</code>  | The time as <code>%H:%M:%S</code> .                                                                                                                                                                                                                                                                                                                                     |
| 44608 | <code>%U</code>  | The week number of the year (Sunday as the first day of the week) as a decimal number [00,53]; leading zeros are permitted but not required.                                                                                                                                                                                                                            |
| 44609 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44610 | <code>%w</code>  | The weekday as a decimal number [0,6], with 0 representing Sunday; leading zeros are permitted but not required.                                                                                                                                                                                                                                                        |
| 44611 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44612 | <code>%W</code>  | The week number of the year (Monday as the first day of the week) as a decimal number [00,53]; leading zeros are permitted but not required.                                                                                                                                                                                                                            |
| 44613 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44614 | <code>%x</code>  | The date, using the locale's date format.                                                                                                                                                                                                                                                                                                                               |
| 44615 | <code>%X</code>  | The time, using the locale's time format.                                                                                                                                                                                                                                                                                                                               |
| 44616 | <code>%y</code>  | The year within century. When a century is not otherwise specified, values in the range 69-99 refer to years in the twentieth century (1969 to 1999 inclusive); values in the range 00-68 refer to years in the twenty-first century (2000 to 2068 inclusive); leading zeros are permitted but not required.                                                            |
| 44617 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44618 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44619 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44620 | <code>%Y</code>  | The year, including the century (for example, 1988).                                                                                                                                                                                                                                                                                                                    |
| 44621 | <code>%%</code>  | Replaced by <code>'%'</code> .                                                                                                                                                                                                                                                                                                                                          |
| 44622 |                  | <b>Modified Directives</b>                                                                                                                                                                                                                                                                                                                                              |
| 44623 |                  | Some directives can be modified by the <i>E</i> and <i>O</i> modifier characters to indicate that an alternative format or specification should be used rather than the one normally used by the unmodified directive. If the alternative format or specification does not exist in the current locale, the behavior shall be as if the unmodified directive were used. |
| 44624 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44625 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44626 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44627 | <code>%Ec</code> | The locale's alternative appropriate date and time representation.                                                                                                                                                                                                                                                                                                      |
| 44628 | <code>%EC</code> | The name of the base year (period) in the locale's alternative representation.                                                                                                                                                                                                                                                                                          |
| 44629 | <code>%Ex</code> | The locale's alternative date representation.                                                                                                                                                                                                                                                                                                                           |
| 44630 | <code>%EX</code> | The locale's alternative time representation.                                                                                                                                                                                                                                                                                                                           |
| 44631 | <code>%Ey</code> | The offset from <code>%EC</code> (year only) in the locale's alternative representation.                                                                                                                                                                                                                                                                                |
| 44632 | <code>%EY</code> | The full alternative year representation.                                                                                                                                                                                                                                                                                                                               |
| 44633 | <code>%Od</code> | The day of the month using the locale's alternative numeric symbols; leading zeros are permitted but not required.                                                                                                                                                                                                                                                      |
| 44634 |                  |                                                                                                                                                                                                                                                                                                                                                                         |
| 44635 | <code>%Oe</code> | The same as <code>%Od</code> .                                                                                                                                                                                                                                                                                                                                          |
| 44636 | <code>%OH</code> | The hour (24-hour clock) using the locale's alternative numeric symbols.                                                                                                                                                                                                                                                                                                |
| 44637 | <code>%OI</code> | The hour (12-hour clock) using the locale's alternative numeric symbols.                                                                                                                                                                                                                                                                                                |
| 44638 | <code>%Om</code> | The month using the locale's alternative numeric symbols.                                                                                                                                                                                                                                                                                                               |
| 44639 | <code>%OM</code> | The minutes using the locale's alternative numeric symbols.                                                                                                                                                                                                                                                                                                             |
| 44640 | <code>%OS</code> | The seconds using the locale's alternative numeric symbols.                                                                                                                                                                                                                                                                                                             |
| 44641 | <code>%OU</code> | The week number of the year (Sunday as the first day of the week) using the locale's alternative numeric symbols.                                                                                                                                                                                                                                                       |
| 44642 |                  |                                                                                                                                                                                                                                                                                                                                                                         |

- 44643 %Ow The number of the weekday (Sunday=0) using the locale's alternative numeric symbols.
- 44644 %OW The week number of the year (Monday as the first day of the week) using the locale's  
44645 alternative numeric symbols.
- 44646 %Oy The year (offset from %C) using the locale's alternative numeric symbols.
- 44647 A directive composed of white-space characters is executed by scanning input up to the first  
44648 character that is not white-space (which remains unscanned), or until no more characters can be  
44649 scanned.
- 44650 A directive that is an ordinary character is executed by scanning the next character from the  
44651 buffer. If the character scanned from the buffer differs from the one comprising the directive, the  
44652 directive fails, and the differing and subsequent characters remain unscanned.
- 44653 A series of directives composed of %n, %t, white-space characters, or any combination is  
44654 executed by scanning up to the first character that is not white space (which remains  
44655 unscanned), or until no more characters can be scanned.
- 44656 Any other conversion specification is executed by scanning characters until a character matching  
44657 the next directive is scanned, or until no more characters can be scanned. These characters,  
44658 except the one matching the next directive, are then compared to the locale values associated  
44659 with the conversion specifier. If a match is found, values for the appropriate **tm** structure  
44660 members are set to values corresponding to the locale information. Case is ignored when  
44661 matching items in *buf* such as month or weekday names. If no match is found, *strptime()* fails  
44662 and no more characters are scanned.
- 44663 **RETURN VALUE**
- 44664 Upon successful completion, *strptime()* shall return a pointer to the character following the last  
44665 character parsed. Otherwise, a null pointer shall be returned.
- 44666 **ERRORS**
- 44667 No errors are defined.
- 44668 **EXAMPLES**
- 44669 None.
- 44670 **APPLICATION USAGE**
- 44671 Several "same as" formats and the special processing of white-space characters are provided in  
44672 order to ease the use of identical *format* strings for *strftime()* and *strptime()*.
- 44673 Applications should use %Y (4-digit years) in preference to %y (2-digit years).
- 44674 It is unspecified whether multiple calls to *strptime()* using the same **tm** structure will update the  
44675 current contents of the structure or overwrite all contents of the structure. Portable applications  
44676 should make a single call to *strptime()* with a format and all data needed to completely specify  
44677 the date and time being converted.
- 44678 **RATIONALE**
- 44679 None.
- 44680 **FUTURE DIRECTIONS**
- 44681 The *strptime()* function is expected to be mandatory in the next version of this volume of  
44682 IEEE Std. 1003.1-200x.
- 44683 **SEE ALSO**
- 44684 *scanf()*, *strftime()*, *time()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

44685 **CHANGE HISTORY**

44686 First released in Issue 4.

44687 **Issue 5**

44688 Moved from ENHANCED I18N to BASE.

44689 The [ENOSYS] error is removed.

44690 The exact meaning of the %y and %Oy specifiers are clarified in the DESCRIPTION.

44691 **Issue 6**

44692 The Open Group corrigenda item U033/5 has been applied. The %r specifier description is reworded.

44694 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

44695 The **restrict** keyword is added to the *strptime()* prototype for alignment with the |  
44696 ISO/IEC 9899:1999 standard. |

44697 The Open Group corrigenda item U047/2 has been applied. |

44698 **NAME**

44699           strchr — string scanning operation

44700 **SYNOPSIS**

44701           #include &lt;string.h&gt;

44702           char \*strchr(const char \*s, int c);

44703 **DESCRIPTION**44704 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
44705 conflict between the requirements described here and the ISO C standard is unintentional. This  
44706 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.44707 **CX**       The *strchr()* function shall locate the last occurrence of *c* (converted to an **unsigned char**) in the  
44708 string pointed to by *s*. The terminating null byte is considered to be part of the string.44709 **RETURN VALUE**44710           Upon successful completion, *strchr()* shall return a pointer to the byte or a null pointer if *c* does  
44711 not occur in the string.44712 **ERRORS**

44713           No errors are defined.

44714 **EXAMPLES**44715           **Finding the Base Name of a File**44716           The following example uses *strchr()* to get a pointer to the base name of a file. The *strchr()*  
44717 function searches backwards through the name of the file to find the last '/' character in *name*.  
44718 This pointer (plus one) will point to the base name of the file.

44719           #include &lt;string.h&gt;

44720           ...

44721           const char \*name;

44722           char \*basename;

44723           ...

44724           basename = strchr(name, '/') + 1;

44725           ...

44726 **APPLICATION USAGE**

44727           None.

44728 **RATIONALE**

44729           None.

44730 **FUTURE DIRECTIONS**

44731           None.

44732 **SEE ALSO**44733           *strchr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>44734 **CHANGE HISTORY**

44735           First released in Issue 1. Derived from Issue 1 of the SVID.

44736 **Issue 4**44737           The **DESCRIPTION** and **RETURN VALUE** sections are changed to make it clear that the function  
44738 works in units of bytes rather than (possibly multi-byte) characters.

44739           The following change is incorporated for alignment with the ISO C standard:

44740

- The type of argument *s* is changed from **char\*** to **const char\***.

44741 **NAME**

44742           strspn — get length of a substring

44743 **SYNOPSIS**

44744           #include &lt;string.h&gt;

44745           size\_t strspn(const char \*s1, const char \*s2);

44746 **DESCRIPTION**

44747 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
44748 conflict between the requirements described here and the ISO C standard is unintentional. This  
44749 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44750       The *strspn()* function shall compute the length of the maximum initial segment of the string  
44751 pointed to by *s1* which consists entirely of bytes from the string pointed to by *s2*.

44752 **RETURN VALUE**

44753       The *strspn()* function shall return the length of *s1*; no return value is reserved to indicate an  
44754 error.

44755 **ERRORS**

44756       No errors are defined.

44757 **EXAMPLES**

44758       None.

44759 **APPLICATION USAGE**

44760       None.

44761 **RATIONALE**

44762       None.

44763 **FUTURE DIRECTIONS**

44764       None.

44765 **SEE ALSO**44766       *strcspn()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**string.h**>44767 **CHANGE HISTORY**

44768       First released in Issue 1. Derived from Issue 1 of the SVID.

44769 **Issue 4**

44770       The **DESCRIPTION** is changed to make it clear that the function works in units of bytes rather  
44771 than (possibly multi-byte) characters.

44772       The following change is incorporated for alignment with the ISO C standard:

- 44773       • The type of arguments *s1* and *s2* are changed from **char\*** to **const char\***.

44774 **Issue 5**

44775       The **RETURN VALUE** section is updated to indicate that *strspn()* returns the length of *s*, and not  
44776 *s* itself as was previously stated.

44777 **NAME**

44778            strstr — find a substring

44779 **SYNOPSIS**

44780            #include &lt;string.h&gt;

44781            char \*strstr(const char \*s1, const char \*s2);

44782 **DESCRIPTION**

44783 cx        The functionality described on this reference page is aligned with the ISO C standard. Any  
44784        conflict between the requirements described here and the ISO C standard is unintentional. This  
44785        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44786        The *strstr()* function shall locate the first occurrence in the string pointed to by *s1* of the  
44787        sequence of bytes (excluding the terminating null byte) in the string pointed to by *s2*.

44788 **RETURN VALUE**

44789        Upon successful completion, *strstr()* shall return a pointer to the located string or a null pointer  
44790        if the string is not found.

44791        If *s2* points to a string with zero length, the function shall return *s1*.

44792 **ERRORS**

44793        No errors are defined.

44794 **EXAMPLES**

44795        None.

44796 **APPLICATION USAGE**

44797        None.

44798 **RATIONALE**

44799        None.

44800 **FUTURE DIRECTIONS**

44801        None.

44802 **SEE ALSO**44803        *strchr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>44804 **CHANGE HISTORY**

44805        First released in Issue 3.

44806        Entry included for alignment with the ANSI C standard.

44807 **Issue 4**

44808        The DESCRIPTION is changed to make it clear that the function works in units of bytes rather  
44809        than (possibly multi-byte) characters.

44810        The following change is incorporated for alignment with the ISO C standard:

- 44811        • The type of arguments *s1* and *s2* are changed from **char\*** to **const char\***.

## 44812 NAME

44813 strtod, strtodf, strtold — convert string to a double-precision number

## 44814 SYNOPSIS

44815 #include <stdlib.h>

44816 double strtod(const char \*restrict *nptr*, char \*\*restrict *endptr*);

44817 float strtodf(const char \*restrict *nptr*, char \*\*restrict *endptr*);

44818 long double strtold(const char \*restrict *nptr*, char \*\*restrict *endptr*);

## 44819 DESCRIPTION

44820 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
44821 conflict between the requirements described here and the ISO C standard is unintentional. This  
44822 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44823 These functions shall convert the initial portion of the string pointed to by *nptr* to **double**, **float**,  
44824 and **long double** representation, respectively. First, they decompose the input string into three  
44825 parts:

- 44826 1. An initial, possibly empty, sequence of white-space characters (as specified by *isspace()*)
- 44827 2. A subject sequence interpreted as a floating-point constant or representing infinity or NaN
- 44828 3. A final string of one or more unrecognized characters, including the terminating null byte  
44829 of the input string

44830 Then it attempts to convert the subject sequence to a floating-point number, and returns the  
44831 result.

44832 The expected form of the subject sequence is an optional plus or minus sign, then one of the  
44833 following:

- 44834 • A non-empty sequence of decimal digits optionally containing a radix character, then an  
44835 optional exponent part
- 44836 • A 0x or 0X, then a non-empty sequence of hexadecimal digits optionally containing a radix  
44837 character, then an optional binary exponent part
- 44838 • One of INF or INFINITY, ignoring case
- 44839 • One of NAN or NAN(*n-char-sequence<sub>opt</sub>*), ignoring case in the NAN part, where:

44840 n-char-sequence:

44841 digit

44842 nondigit

44843 n-char-sequence digit

44844 n-char-sequence nondigit

44845 The subject sequence is defined as the longest initial subsequence of the input string, starting  
44846 with the first non-white-space character, that is of the expected form. The subject sequence  
44847 contains no characters if the input string is not of the expected form.

44848 If the subject sequence has the expected form for a floating-point number, the sequence of  
44849 characters starting with the first digit or the decimal-point character (whichever occurs first) is  
44850 interpreted as a floating constant of the C language, except that the radix character is used in  
44851 place of a period, and that if neither an exponent part nor a radix character appears in a decimal  
44852 floating-point number, or if a binary exponent part does not appear in a hexadecimal floating-  
44853 point number, an exponent part of the appropriate type with value zero is assumed to follow the  
44854 last digit in the string. If the subject sequence begins with a minus sign, the sequence is  
44855 interpreted as negated. A character sequence INF or INFINITY is interpreted as an infinity, if

44856 representable in the return type, else like a floating constant that is too large for the range of the  
 44857 return type. A character sequence NAN or NAN(*n-char-sequence<sub>opt</sub>*), is interpreted as a quiet NaN, if  
 44858 supported in the return type, else like a subject sequence part that does not have the expected form; the  
 44859 meaning of the *n-char* sequences is implementation-defined. A pointer to the final string is stored  
 44860 in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

44861 If the subject sequence has the hexadecimal form and FLT\_RADIX is a power of 2, the value  
 44862 resulting from the conversion is correctly rounded.

44863 CX The radix character is defined in the program's locale (category *LC\_NUMERIC*). In the POSIX  
 44864 locale, or in a locale where the radix character is not defined, the radix character defaults to a  
 44865 period ('.').

44866 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be  
 44867 accepted.

44868 If the subject sequence is empty or does not have the expected form, no conversion is performed;  
 44869 the value of *str* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null  
 44870 pointer.

44871 The *strtod()* function shall not change the setting of *errno* if successful.

44872 Because 0 is returned on error and is also a valid return on success, an application wishing to  
 44873 check for error situations should set *errno* to 0, then call *strtod()*, then check *errno*.

#### 44874 RETURN VALUE

44875 Upon successful completion, these functions shall return the converted value. If no conversion  
 44876 could be performed, 0 shall be returned, and *errno* may be set to [EINVAL].

44877 If the correct value is outside the range of representable values, HUGE\_VAL, HUGE\_VALF, or  
 44878 HUGE\_VALL shall be returned (according to the sign of the value), and *errno* shall be set to  
 44879 [ERANGE].

44880 If the correct value would cause an underflow, a value whose magnitude is no greater than the  
 44881 smallest normalized positive number in the return type shall be returned and *errno* set to  
 44882 [ERANGE].

#### 44883 ERRORS

44884 The *strtod()* function shall fail if:

44885 CX [ERANGE] The value to be returned would cause overflow or underflow.

44886 The *strtod()* function may fail if:

44887 CX [EINVAL] No conversion could be performed.

#### 44888 Notes to Reviewers

44889 *This section with side shading will not appear in the final copy. - Ed.*

44890 There is a query outstanding over the question of [EINVAL] being an allowable extension to  
 44891 C99. This error may be removed in a future draft.

44892 **EXAMPLES**

44893 None.

44894 **APPLICATION USAGE**

44895 If the subject sequence has the hexadecimal form and FLT\_RADIX is not a power of 2, the result  
 44896 should be one of the two numbers in the appropriate internal format that are adjacent to the  
 44897 hexadecimal floating source value, with the extra stipulation that the error should have a correct  
 44898 sign for the current rounding direction.

44899 If the subject sequence has the decimal form and at most DECIMAL\_DIG (defined in <float.h>)  
 44900 significant digits, the result should be correctly rounded. If the subject sequence *D* has the  
 44901 decimal form and more than DECIMAL\_DIG significant digits, consider the two bounding,  
 44902 adjacent decimal strings *L* and *U*, both having DECIMAL\_DIG significant digits, such that the  
 44903 values of *L*, *D*, and *U* satisfy " $L \leq D \leq U$ ". The result should be one of the (equal or  
 44904 adjacent) values that would be obtained by correctly rounding *L* and *U* according to the current  
 44905 rounding direction, with the extra stipulation that the error with respect to *D* should have a  
 44906 correct sign for the current rounding direction.

44907 **RATIONALE**

44908 None.

44909 **FUTURE DIRECTIONS**

44910 None.

44911 **SEE ALSO**

44912 *isspace()*, *localeconv()*, *scanf()*, *setlocale()*, *strtol()*, the Base Definitions volume of  
 44913 IEEE Std. 1003.1-200x, <float.h>, <stdlib.h>, the Base Definitions volume of  
 44914 IEEE Std. 1003.1-200x, Chapter 7, Locale

44915 **CHANGE HISTORY**

44916 First released in Issue 1. Derived from Issue 1 of the SVID.

44917 **Issue 4**

44918 The DESCRIPTION is changed to make it clear when the function manipulates bytes and when  
 44919 it manipulates characters.

44920 The [EINVAL] error is added to the ERRORS section and marked as an extension.

44921 The following changes are incorporated for alignment with the ISO C standard:

- 44922 • The function is no longer marked as an extension.
- 44923 • The type of argument *str* is changed from **char\*** to **const char\***.
- 44924 • The name of the second argument is changed from *ptr* to *endptr*.
- 44925 • The precise conditions under which the [ERANGE] error can be set have been defined in the  
 44926 RETURN VALUE section.

44927 **Issue 5**

44928 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

44929 **Issue 6**

44930 Extensions beyond the ISO C standard are now marked.

44931 The following new requirements on POSIX implementations derive from alignment with the  
 44932 Single UNIX Specification:

- 44933 • In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
 44934 added if no conversion could be performed.

- 44935 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:
- 44936 • The *strtod()* function is updated.
  - 44937 • The *strtof()* and *strtold()* functions are added.
  - 44938 • The DESCRIPTION is extensively revised.

44939 **NAME**

44940 strtoimax, strtoumax — convert string to integer type

44941 **SYNOPSIS**

44942 #include <inttypes.h>

44943 intmax\_t strtoimax(const char \*restrict *nptr*, char \*\*restrict *endptr*,  
44944 int *base*);

44945 uintmax\_t strtoumax(const char \*restrict *nptr*, char \*\*restrict *endptr*,  
44946 int *base*);

44947 **DESCRIPTION**

44948 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
44949 conflict between the requirements described here and the ISO C standard is unintentional. This  
44950 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44951 These functions shall be equivalent to the *strtol()*, *strtoll()*, *strtoul()*, and *strtoull()* functions,  
44952 except that the initial portion of the string shall be converted to **intmax\_t** and **uintmax\_t**  
44953 representation, respectively.

44954 **RETURN VALUE**

44955 These functions shall return the converted value, if any.

44956 If no conversion could be performed, zero shall be returned.

44957 If the correct value is outside the range of representable values, {INTMAX\_MAX},  
44958 {INTMAX\_MIN}, or {UINTMAX\_MAX} shall be returned (according to the return type and sign  
44959 of the value, if any), and *errno* shall be set to [ERANGE].

44960 **ERRORS**

44961 These functions shall fail if:

44962 [ERANGE] The value to be returned is not representable.

44963 These functions may fail if:

44964 [EINVAL] The value of *base* is not supported.

44965 **EXAMPLES**

44966 None.

44967 **APPLICATION USAGE**

44968 None.

44969 **RATIONALE**

44970 None.

44971 **FUTURE DIRECTIONS**

44972 None.

44973 **SEE ALSO**

44974 *strtol()*, *strtoul()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <inttypes.h>

44975 **CHANGE HISTORY**

44976 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

## 44977 NAME

44978 strtok, strtok\_r — split string into tokens

## 44979 SYNOPSIS

44980 #include &lt;string.h&gt;

44981 char \*strtok(char \*restrict s1, const char \*restrict s2);

44982 TSF char \*strtok\_r(char \*restrict s, const char \*restrict sep,

44983 char \*\*restrict lasts);

44984

## 44985 DESCRIPTION

44986 CX For *strtok()*: The functionality described on this reference page is aligned with the ISO C standard. Any conflict between the requirements described here and the ISO C standard is unintentional. This volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

44989 A sequence of calls to *strtok()* breaks the string pointed to by *s1* into a sequence of tokens, each of which is delimited by a byte from the string pointed to by *s2*. The first call in the sequence has *s1* as its first argument, and is followed by calls with a null pointer as their first argument. The separator string pointed to by *s2* may be different from call to call.

44993 The first call in the sequence searches the string pointed to by *s1* for the first byte that is *not* contained in the current separator string pointed to by *s2*. If no such byte is found, then there are no tokens in the string pointed to by *s1* and *strtok()* returns a null pointer. If such a byte is found, it is the start of the first token.

44997 The *strtok()* function then searches from there for a byte that *is* contained in the current separator string. If no such byte is found, the current token extends to the end of the string pointed to by *s1*, and subsequent searches for a token shall return a null pointer. If such a byte is found, it is overwritten by a null byte, which terminates the current token. The *strtok()* function saves a pointer to the following byte, from which the next search for a token shall start.

45002 Each subsequent call, with a null pointer as the value of the first argument, starts searching from the saved pointer and behaves as described above.

45004 The implementation shall behave as if no function defined in this volume of IEEE Std. 1003.1-200x calls *strtok()*.

45006 CX The *strtok()* function need not be reentrant. A function that is not required to be reentrant is not required to be thread-safe.

45008 TSF The *strtok\_r()* function considers the null-terminated string *s* as a sequence of zero or more text tokens separated by spans of one or more characters from the separator string *sep*. The argument *lasts* points to a user-provided pointer which points to stored information necessary for *strtok\_r()* to continue scanning the same string.

45012 In the first call to *strtok\_r()*, *s* points to a null-terminated string, *sep* to a null-terminated string of separator characters, and the value pointed to by *lasts* is ignored. The *strtok\_r()* function returns a pointer to the first character of the first token, writes a null character into *s* immediately following the returned token, and updates the pointer to which *lasts* points.

45016 In subsequent calls, *s* is a NULL pointer and *lasts* shall be unchanged from the previous call so that subsequent calls shall move through the string *s*, returning successive tokens until no tokens remain. The separator string *sep* may be different from call to call. When no token remains in *s*, a NULL pointer is returned.

## 45020 RETURN VALUE

45021 Upon successful completion, *strtok()* shall return a pointer to the first byte of a token. Otherwise,  
45022 if there is no token, *strtok()* shall return a null pointer.

45023 TSF The *strtok\_r()* function shall return a pointer to the token found, or a NULL pointer when no  
45024 token is found.

## 45025 ERRORS

45026 No errors are defined.

## 45027 EXAMPLES

## 45028 Searching for Word Separators

45029 The following example searches for tokens separated by space characters.

```
45030 #include <string.h>
45031 ...
45032 char *token;
45033 char *line = "LINE TO BE SEPARATED";
45034 char *search = " ";

45035 /* Token will point to "LINE". */
45036 token = strtok(line, search);

45037 /* Token will point to "TO". */
45038 token = strtok(NULL, search);
```

## 45039 Breaking a Line

45040 The following example uses *strtok()* to break a line into two character strings separated by any  
45041 combination of <space>s, <tab>s, or <newline>s.

```
45042 #include <string.h>
45043 ...
45044 struct element {
45045 char *key;
45046 char *data;
45047 };
45048 ...
45049 char line[LINE_MAX];
45050 char *key, *data;
45051 ...
45052 key = strtok(line, " \n");
45053 data = strtok(NULL, " \n");
45054 ...
```

## 45055 APPLICATION USAGE

45056 The *strtok\_r()* function is thread-safe and stores its state in a user-supplied buffer instead of  
45057 possibly using a static data area that may be overwritten by an unrelated call from another  
45058 thread.

## 45059 RATIONALE

45060 The *strtok()* function searches for a separator string within a larger string. It returns a pointer to  
45061 the last substring between separator strings. This function uses static storage to keep track of  
45062 the current string position between calls. The new function, *strtok\_r()*, takes an additional  
45063 argument, *lasts*, to keep track of the current position in the string.

45064 **FUTURE DIRECTIONS**

45065 None.

45066 **SEE ALSO**

45067 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;string.h&gt;

45068 **CHANGE HISTORY**

45069 First released in Issue 1. Derived from Issue 1 of the SVID.

45070 **Issue 4**45071 The DESCRIPTION is changed to make it clear that the function manipulates bytes rather than  
45072 (possibly multi-byte) characters.

45073 The following changes are incorporated for alignment with the ISO C standard:

- 45074
- The function is no longer marked as an extension.
  - The type of argument *s2* is changed from **char\*** to **const char\***.
- 45075

45076 **Issue 5**45077 The *strtok\_r()* function is included for alignment with the POSIX Threads Extension.45078 A note indicating that the *strtok()* function need not be reentrant is added to the DESCRIPTION.45079 **Issue 6**

45080 Extensions beyond the ISO C standard are now marked.

45081 The *strtok\_r()* function is marked as part of the Thread-Safe Functions option.

45082 In the DESCRIPTION, the note about reentrancy is expanded to cover thread-safety.

45083 The APPLICATION USAGE section is updated to include a note on the thread-safe function and  
45084 its avoidance of possibly using a static data area.45085 The **restrict** keyword is added to the *strtok()* and *strtok\_r()* prototypes for alignment with the  
45086 ISO/IEC 9899:1999 standard.

## 45087 NAME

45088 strtol, strtoll — convert string to a long integer

## 45089 SYNOPSIS

45090 #include &lt;stdlib.h&gt;

45091 long strtol(const char \*restrict *str*, char \*\*restrict *endptr*, int *base*);45092 long long strtoll(const char \*restrict *str*, char \*\*restrict *endptr*,45093 int *base*)

## 45094 DESCRIPTION

45095 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 45096 conflict between the requirements described here and the ISO C standard is unintentional. This  
 45097 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

45098 These functions shall convert the initial portion of the string pointed to by *str* to a type **long** and  
 45099 **long long** representation, respectively. First, they decompose the input string into three parts:

- 45100 1. An initial, possibly empty, sequence of white-space characters (as specified by *isspace()*)
- 45101 2. A subject sequence interpreted as an integer represented in some radix determined by the  
 45102 value of *base*
- 45103 3. A final string of one or more unrecognized characters, including the terminating null byte  
 45104 of the input string.

45105 Then it attempts to convert the subject sequence to an integer, and returns the result.

45106 If the value of *base* is 0, the expected form of the subject sequence is that of a decimal constant,  
 45107 octal constant, or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A  
 45108 decimal constant begins with a non-zero digit, and consists of a sequence of decimal digits. An  
 45109 octal constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to  
 45110 '7' only. A hexadecimal constant consists of the prefix "0x" or "0X" followed by a sequence of  
 45111 the decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.

45112 If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence  
 45113 of letters and digits representing an integer with the radix specified by *base*, optionally preceded  
 45114 by a '+' or '-' sign. The letters from 'a' (or 'A') to 'z' (or 'Z') inclusive are ascribed the  
 45115 values 10 to 35; only letters whose ascribed values are less than that of *base* are permitted. If the  
 45116 value of *base* is 16, the characters "0x" or "0X" may optionally precede the sequence of letters  
 45117 and digits, following the sign if present.

45118 The subject sequence is defined as the longest initial subsequence of the input string, starting  
 45119 with the first non-white-space character that is of the expected form. The subject sequence  
 45120 contains no characters if the input string is empty or consists entirely of white-space characters,  
 45121 or if the first non-white-space character is other than a sign or a permissible letter or digit.

45122 If the subject sequence has the expected form and the value of *base* is 0, the sequence of  
 45123 characters starting with the first digit is interpreted as an integer constant. If the subject  
 45124 sequence has the expected form and the value of *base* is between 2 and 36, it is used as the base  
 45125 for conversion, ascribing to each letter its value as given above. If the subject sequence begins  
 45126 with a minus sign, the value resulting from the conversion is negated. A pointer to the final  
 45127 string is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

45128 cx In other than the C or POSIX locales, other implementation-defined subject sequences may be  
 45129 accepted.

45130 If the subject sequence is empty or does not have the expected form, no conversion is performed;  
 45131 the value of *str* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null

45132 pointer.

45133 The *strtol()* function shall not change the setting of *errno* if successful.

45134 Because 0, {LONG\_MIN} or {LLONG\_MIN}, and {LONG\_MAX} or {LLONG\_MAX} are returned  
45135 on error and are also valid returns on success, an application wishing to check for error  
45136 situations should set *errno* to 0, then call *strtol()*, then check *errno*.

#### 45137 RETURN VALUE

45138 Upon successful completion, these functions shall return the converted value, if any. If no  
45139 CX conversion could be performed, 0 shall be returned and *errno* may be set to [EINVAL].

45140 If the correct value is outside the range of representable values, {LONG\_MIN}, {LONG\_MAX},  
45141 {LLONG\_MIN} or {LLONG\_MAX} shall be returned (according to the sign of the value), and  
45142 *errno* set to [ERANGE].

#### 45143 ERRORS

45144 The *strtol()* function shall fail if:

45145 [ERANGE] The value to be returned is not representable.

45146 The *strtol()* function may fail if:

45147 CX [EINVAL] The value of *base* is not supported.

#### 45148 EXAMPLES

45149 None.

#### 45150 APPLICATION USAGE

45151 None.

#### 45152 RATIONALE

45153 None.

#### 45154 FUTURE DIRECTIONS

45155 None.

#### 45156 SEE ALSO

45157 *isalpha()*, *scanf()*, *strtod()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>

#### 45158 CHANGE HISTORY

45159 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 45160 Issue 4

45161 The DESCRIPTION is changed to make it clear when the function manipulates bytes and when  
45162 it manipulates characters.

45163 In the RETURN VALUE section, text indicating that *errno* is set when 0 is returned is marked as  
45164 an extension.

45165 The ERRORS section is updated in line with the RETURN VALUE section.

45166 The following changes are incorporated for alignment with the ISO C standard:

- 45167 • The function is no longer marked as an extension.
- 45168 • The type of argument *str* is changed from **char\*** to **const char\***.
- 45169 • The name of the second argument is changed from *ptr* to *endptr*.
- 45170 • The DESCRIPTION is changed to indicate permitted forms of the subject sequence when *base*  
45171 is 0.

- 45172           • The RETURN VALUE section is changed to indicate that {LONG\_MAX} or {LONG\_MIN} is
- 45173            returned if the converted value is too large or too small.
  
- 45174 **Issue 5**
- 45175           The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.
  
- 45176 **Issue 6**
- 45177           Extensions beyond the ISO C standard are now marked.
  
- 45178           The following new requirements on POSIX implementations derive from alignment with the
- 45179           Single UNIX Specification:
  
- 45180           • In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is
- 45181            added if no conversion could be performed.
  
- 45182           The following changes are made for alignment with the ISO/IEC 9899:1999 standard:
  
- 45183           • The *strtol()* prototype is updated.
  
- 45184           • The *strtoll()* function is added.

## 45185 NAME

45186 strtoul, strtoull — convert string to an unsigned long

## 45187 SYNOPSIS

45188 #include &lt;stdlib.h&gt;

45189 long strtoul(const char \*restrict *str*, char \*\*restrict *endptr*, int *base*);45190 long long strtoull(const char\*restrict *str*, char \*\*restrict *endptr*,45191 int *base*);

## 45192 DESCRIPTION

45193 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 45194 conflict between the requirements described here and the ISO C standard is unintentional. This  
 45195 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

45196 These functions shall convert the initial portion of the string pointed to by *str* to a type **long** and  
 45197 **long long** representation, respectively. First, they decompose the input string into three parts:

- 45198 1. An initial, possibly empty, sequence of white-space characters (as specified by *isspace()*)
- 45199 2. A subject sequence interpreted as an integer represented in some radix determined by the  
 45200 value of *base*
- 45201 3. A final string of one or more unrecognized characters, including the terminating null byte  
 45202 of the input string

45203 Then it attempts to convert the subject sequence to an unsigned integer, and returns the result.

45204 If the value of *base* is 0, the expected form of the subject sequence is that of a decimal constant,  
 45205 octal constant, or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A  
 45206 decimal constant begins with a non-zero digit, and consists of a sequence of decimal digits. An  
 45207 octal constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to  
 45208 '7' only. A hexadecimal constant consists of the prefix "0x" or "0X" followed by a sequence of  
 45209 the decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.

45210 If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence  
 45211 of letters and digits representing an integer with the radix specified by *base*, optionally preceded  
 45212 by a '+' or '-' sign. The letters from 'a' (or 'A') to 'z' (or 'Z') inclusive are ascribed the  
 45213 values 10 to 35; only letters whose ascribed values are less than that of *base* are permitted. If the  
 45214 value of *base* is 16, the characters "0x" or "0X" may optionally precede the sequence of letters  
 45215 and digits, following the sign if present.

45216 The subject sequence is defined as the longest initial subsequence of the input string, starting  
 45217 with the first non-white-space character that is of the expected form. The subject sequence  
 45218 contains no characters if the input string is empty or consists entirely of white-space characters,  
 45219 or if the first non-white-space character is other than a sign or a permissible letter or digit.

45220 If the subject sequence has the expected form and the value of *base* is 0, the sequence of  
 45221 characters starting with the first digit is interpreted as an integer constant. If the subject  
 45222 sequence has the expected form and the value of *base* is between 2 and 36, it is used as the base  
 45223 for conversion, ascribing to each letter its value as given above. If the subject sequence begins  
 45224 with a minus sign, the value resulting from the conversion is negated. A pointer to the final  
 45225 string is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

45226 cx In other than the C or POSIX locales, other implementation-defined subject sequences may be  
 45227 accepted.

45228 If the subject sequence is empty or does not have the expected form, no conversion is performed;  
 45229 the value of *str* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null

45230 pointer.

45231 The *strtol()* function shall not change the setting of *errno* if successful.

45232 Because 0, {ULONG\_MAX}, and {ULLONG\_MAX} are returned on error and are also valid  
 45233 returns on success, an application wishing to check for error situations should set *errno* to 0, then  
 45234 call *strtol()*, then check *errno*.

#### 45235 RETURN VALUE

45236 Upon successful completion, these functions shall return the converted value, if any. If no  
 45237 CX conversion could be performed, 0 shall be returned and *errno* may be set to [EINVAL]. If the  
 45238 correct value is outside the range of representable values, {ULONG\_MAX} or {ULLONG\_MAX}  
 45239 shall be returned and *errno* set to [ERANGE].

#### 45240 ERRORS

45241 The *strtol()* function shall fail if:

45242 CX [EINVAL] The value of *base* is not supported.

45243 [ERANGE] The value to be returned is not representable.

45244 The *strtol()* function may fail if:

45245 CX [EINVAL] No conversion could be performed.

#### 45246 EXAMPLES

45247 None.

#### 45248 APPLICATION USAGE

45249 None.

#### 45250 RATIONALE

45251 None.

#### 45252 FUTURE DIRECTIONS

45253 None.

#### 45254 SEE ALSO

45255 *isalpha()*, *scanf()*, *strtod()*, *strtol()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
 45256 <stdlib.h>

#### 45257 CHANGE HISTORY

45258 First released in Issue 4. Derived from the ANSI C standard.

#### 45259 Issue 5

45260 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

#### 45261 Issue 6

45262 Extensions beyond the ISO C standard are now marked.

45263 The following new requirements on POSIX implementations derive from alignment with the  
 45264 Single UNIX Specification:

- 45265 • The [EINVAL] error condition is added for when the value of *base* is not supported.

45266 In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
 45267 added if no conversion could be performed.

45268 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

- 45269 • The *strtol()* prototype is updated.

45270

- The *strtoull()* function is added.

45271 **NAME**

45272 strxfrm — string transformation

45273 **SYNOPSIS**

45274 #include &lt;string.h&gt;

45275 size\_t strxfrm(char \*restrict *s1*, const char \*restrict *s2*, size\_t *n*);45276 **DESCRIPTION**

45277 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 45278 conflict between the requirements described here and the ISO C standard is unintentional. This  
 45279 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

45280 The *strxfrm()* function shall transform the string pointed to by *s2* and place the resulting string  
 45281 into the array pointed to by *s1*. The transformation is such that if *strcmp()* is applied to two  
 45282 transformed strings, it returns a value greater than, equal to, or less than 0, corresponding to the  
 45283 result of *strcoll()* applied to the same two original strings. No more than *n* bytes are placed into  
 45284 the resulting array pointed to by *s1*, including the terminating null byte. If *n* is 0, *s1* is permitted  
 45285 to be a null pointer. If copying takes place between objects that overlap, the behavior is  
 45286 undefined.

45287 CX The *strxfrm()* function shall not change the setting of *errno* if successful.

45288 Because no return value is reserved to indicate an error, an application wishing to check for error  
 45289 situations should set *errno* to 0, then call *strcoll()*, then check *errno*.

45290 **RETURN VALUE**

45291 Upon successful completion, *strxfrm()* shall return the length of the transformed string (not  
 45292 including the terminating null byte). If the value returned is *n* or more, the contents of the array  
 45293 pointed to by *s1* are indeterminate.

45294 CX On error, *strxfrm()* may set *errno* but no return value is reserved to indicate an error.

45295 **ERRORS**

45296 The *strxfrm()* function may fail if:

45297 CX [EINVAL] The string pointed to by the *s2* argument contains characters outside the  
 45298 domain of the collating sequence.

45299 **EXAMPLES**

45300 None.

45301 **APPLICATION USAGE**

45302 The transformation function is such that two transformed strings can be ordered by *strcmp()* as  
 45303 appropriate to collating sequence information in the program's locale (category *LC\_COLLATE*).

45304 The fact that when *n* is 0 *s1* is permitted to be a null pointer is useful to determine the size of the  
 45305 *s1* array prior to making the transformation.

45306 **RATIONALE**

45307 None.

45308 **FUTURE DIRECTIONS**

45309 None.

45310 **SEE ALSO**

45311 *strcmp()*, *strcoll()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <string.h>

45312 **CHANGE HISTORY**

45313 First released in Issue 3.

45314 Entry included for alignment with the ISO C standard.

45315 **Issue 4**45316 The DESCRIPTION is changed to make it clear when the function manipulates byte values and  
45317 when it manipulates characters.45318 The sentence describing error returns in the RETURN VALUE section is marked as an extension,  
45319 as is the [EINVAL] error.

45320 The APPLICATION USAGE section is expanded.

45321 The following changes are incorporated for alignment with the ISO C standard:

- 45322
- The function is no longer marked as an extension.
- 45323
- The type of argument *s2* is changed from **char\*** to **const char\***.

45324 **Issue 5**45325 The DESCRIPTION is updated to indicate that *errno* does not change if the function is  
45326 successful.45327 **Issue 6**

45328 Extensions beyond the ISO C standard are now marked.

45329 The following new requirements on POSIX implementations derive from alignment with the  
45330 Single UNIX Specification:

- 45331
- In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
45332 added if no conversion could be performed.

45333 The *strxfrm()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard. |

## 45334 NAME

45335 swab — swap bytes

## 45336 SYNOPSIS

45337 XSI #include &lt;unistd.h&gt;

45338 void swab(const void \*restrict src, void \*restrict dest,  
45339 ssize\_t nbytes);

45340

## 45341 DESCRIPTION

45342 The *swab()* function shall copy *nbytes* bytes, which are pointed to by *src*, to the object pointed to  
45343 by *dest*, exchanging adjacent bytes. The *nbytes* argument should be even. If *nbytes* is odd, *swab()*  
45344 copies and exchanges *nbytes*–1 bytes and the disposition of the last byte is unspecified. If  
45345 copying takes place between objects that overlap, the behavior is undefined. If *nbytes* is  
45346 negative, *swab()* does nothing.

## 45347 RETURN VALUE

45348 None.

## 45349 ERRORS

45350 No errors are defined.

## 45351 EXAMPLES

45352 None.

## 45353 APPLICATION USAGE

45354 None.

## 45355 RATIONALE

45356 None.

## 45357 FUTURE DIRECTIONS

45358 None.

## 45359 SEE ALSO

45360 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;unistd.h&gt;

## 45361 CHANGE HISTORY

45362 First released in Issue 1. Derived from Issue 1 of the SVID.

## 45363 Issue 4

45364 The &lt;unistd.h&gt; header is added to the SYNOPSIS section.

45365 The type of argument *src* is changed from **char\*** to **const void\***, *dest* is changed from **char\*** to  
45366 **void\***, and *nbytes* is changed from **int** to **ssize\_t**.

45367 The DESCRIPTION is changed as follows:

- 45368 • States explicitly that copying between overlapping objects results in undefined behavior.
- 45369 • Takes account of the type change to *nbyte*; that is, previously it was defined as **int** and could  
45370 be positive or negative, whereas now it is defined as an **unsigned** type.
- 45371 • A statement about overlapping objects is added.

45372 The APPLICATION USAGE section is removed.

45373 **Issue 6**

45374 The **restrict** keyword is added to the *swab()* prototype for alignment with the  
45375 ISO/IEC 9899:1999 standard.

45376 **NAME**

45377 swapcontext — swap user context

45378 **SYNOPSIS**

45379 XSI `#include <ucontext.h>`

45380 `int swapcontext(ucontext_t *restrict oucp,`  
45381 `const ucontext_t *restrict ucp);`

45382

45383 **DESCRIPTION**

45384 Refer to *makecontext()*.

45385 **NAME**

45386 swprintf — print formatted wide-character output

45387 **SYNOPSIS**

45388 #include &lt;stdio.h&gt;

45389 #include &lt;wchar.h&gt;

45390 int swprintf(wchar\_t \*ws, size\_t n, const wchar\_t \*format, ...);

45391 **DESCRIPTION**45392 Refer to *fwprintf()*.

45393 **NAME**

45394       swscanf — convert formatted wide-character input

45395 **SYNOPSIS**

45396       #include <stdio.h>

45397       #include <wchar.h>

45398       int swscanf(const wchar\_t \*ws, const wchar\_t \*format, ... );

45399 **DESCRIPTION**

45400       Refer to *fwscanf()*.

45401 **NAME**

45402            symlink — make symbolic link to a file

45403 **SYNOPSIS**

45404            #include &lt;unistd.h&gt;

45405            int symlink(const char \*path1, const char \*path2);

45406 **DESCRIPTION**

45407            The *symlink()* function shall create a symbolic link called *path2* that contains the string pointed  
 45408            to by *path1* (*path2* is the name of the symbolic link created, *path1* is the string contained in the  
 45409            symbolic link).

45410            The string pointed to by *path1* shall be treated only as a character string and shall not be  
 45411            validated as a path name.

45412            If the *symlink()* function fails for any reason other than [EIO], any file named by *path2* shall be  
 45413            unaffected.

45414 **RETURN VALUE**

45415            Upon successful completion, *symlink()* shall return 0; otherwise, it shall return -1 and set *errno* to  
 45416            indicate the error.

45417 **ERRORS**45418            The *symlink()* function shall fail if:

45419            [EACCES]            Write permission is denied in the directory where the symbolic link is being  
 45420            created, or search permission is denied for a component of the path prefix of  
 45421            *path2*.

45422            [EEXIST]            The *path2* argument names an existing file or symbolic link.

45423            [EIO]                An I/O error occurs while reading from or writing to the file system.

45424            [ELOOP]            A loop exists in symbolic links encountered during resolution of the *path2*  
 45425            argument.

45426            [ENAMETOOLONG]

45427                The length of the *path2* argument exceeds {PATH\_MAX} or a path name  
 45428                component is longer than {NAME\_MAX} or the length of the *path1* argument  
 45429                is longer than {SYMLINK\_MAX}.

45430            [ENOENT]            A component of *path2* does not name an existing file or *path2* is an empty  
 45431            string.

45432            [ENOSPC]            The directory in which the entry for the new symbolic link is being placed  
 45433            cannot be extended because no space is left on the file system containing the  
 45434            directory, or the new symbolic link cannot be created because no space is left  
 45435            on the file system which shall contain the link, or the file system is out of file-  
 45436            allocation resources.

45437            [ENOTDIR]           A component of the path prefix of *path2* is not a directory.

45438            [EROFS]             The new symbolic link would reside on a read-only file system.

45439            The *symlink()* function may fail if:

45440            [ELOOP]            More than {SYMLOOP\_MAX} symbolic links were encountered during  
 45441            resolution of the *path2* argument.

45442            [ENAMETOOLONG]

45443                As a result of encountering a symbolic link in resolution of the *path2*

45444 argument, the length of the substituted path name string exceeded  
45445 {PATH\_MAX} bytes (including the terminating null byte), or the length of the  
45446 string pointed to by *path1* exceeded {SYMLINK\_MAX}.

**45447 EXAMPLES**

45448 None.

**45449 APPLICATION USAGE**

45450 Like a hard link, a symbolic link allows a file to have multiple logical names. The presence of a  
45451 hard link guarantees the existence of a file, even after the original name has been removed. A  
45452 symbolic link provides no such assurance; in fact, the file named by the *path1* argument need not  
45453 exist when the link is created. A symbolic link can cross file system boundaries.

45454 Normal permission checks are made on each component of the symbolic link path name during  
45455 its resolution.

**45456 RATIONALE**

45457 Since IEEE Std. 1003.1-200x does not require any association of file times with symbolic links,  
45458 there is no requirement that file times be updated by *symlink()*.

**45459 FUTURE DIRECTIONS**

45460 None.

**45461 SEE ALSO**

45462 *lchown()*, *link()*, *lstat()*, *open()*, *readlink()*, *unlink()*, the Base Definitions volume of  
45463 IEEE Std. 1003.1-200x, <**unistd.h**>

**45464 CHANGE HISTORY**

45465 First released in Issue 4, Version 2.

**45466 Issue 5**

45467 Moved from X/OPEN UNIX extension to BASE.

**45468 Issue 6**

45469 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 45470 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.  
45471 This is since behavior may vary from one file system to another.

45472 The following changes were made to align with the IEEE P1003.1a draft standard:

- 45473 • The DESCRIPTION text is updated.
- 45474 • The [ELOOP] optional error condition is added.

45475 **NAME**

45476 sync — schedule file system updates

45477 **SYNOPSIS**

45478 XSI #include &lt;unistd.h&gt;

45479 void sync(void);

45480

45481 **DESCRIPTION**45482 The *sync()* function shall cause all information in memory that updates file systems to be  
45483 scheduled for writing out to all file systems.45484 The writing, although scheduled, is not necessarily complete upon return from *sync()*.45485 **RETURN VALUE**45486 The *sync()* function shall return no value.45487 **ERRORS**

45488 No errors are defined.

45489 **EXAMPLES**

45490 None.

45491 **APPLICATION USAGE**

45492 None.

45493 **RATIONALE**

45494 None.

45495 **FUTURE DIRECTIONS**

45496 None.

45497 **SEE ALSO**45498 *fsync()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>45499 **CHANGE HISTORY**

45500 First released in Issue 4, Version 2.

45501 **Issue 5**

45502 Moved from X/OPEN UNIX extension to BASE.

45503 **NAME**

45504 sysconf — get configurable system variables

45505 **SYNOPSIS**

45506 #include <unistd.h>

45507 long sysconf(int name);

45508 **DESCRIPTION**

45509 The *sysconf()* function provides a method for the application to determine the current value of a  
 45510 configurable system limit or option (*variable*). Support for some system variables is dependent  
 45511 on implementation options (as indicated by the margin codes in the following table). Where an  
 45512 implementation option is not supported, the variable need not be supported.

45513 The *name* argument represents the system variable to be queried. The following table lists the  
 45514 minimal set of system variables from <limits.h> or <unistd.h> that can be returned by *sysconf()*,  
 45515 and the symbolic constants, defined in <unistd.h> that are the corresponding values used for  
 45516 *name*. Support for some configuration variables is dependent on implementation options (see  
 45517 shading and margin codes in the table below). Where an implementation option is not  
 45518 supported, the variable need not be supported.

45519

45520

45521 AIO

45522

45523

45524

45525 XSI

45526

45527

45528

45529

45530

45531

45532

45533 XSI

45534

45535 XSI

45536

45537

45538

45539 TSF

45540

45541

45542

45543 MSG

45544

45545

45546 ADV

45547 BAR

45548 AIO

| Variable                                                                    | Value of Name          |
|-----------------------------------------------------------------------------|------------------------|
| {AIO_LISTIO_MAX}                                                            | _SC_AIO_LISTIO_MAX     |
| {AIO_MAX}                                                                   | _SC_AIO_MAX            |
| {AIO_PRIO_DELTA_MAX}                                                        | _SC_AIO_PRIO_DELTA_MAX |
| {ARG_MAX}                                                                   | _SC_ARG_MAX            |
| {ATEXIT_MAX}                                                                | _SC_ATEXIT_MAX         |
| {BC_BASE_MAX}                                                               | _SC_BC_BASE_MAX        |
| {BC_DIM_MAX}                                                                | _SC_BC_DIM_MAX         |
| {BC_SCALE_MAX}                                                              | _SC_BC_SCALE_MAX       |
| {BC_STRING_MAX}                                                             | _SC_BC_STRING_MAX      |
| {CHILD_MAX}                                                                 | _SC_CHILD_MAX          |
| Clock ticks/second                                                          | _SC_CLK_TCK            |
| {COLL_WEIGHTS_MAX}                                                          | _SC_COLL_WEIGHTS_MAX   |
| {DELAYTIMER_MAX}                                                            | _SC_DELAYTIMER_MAX     |
| {EXPR_NEST_MAX}                                                             | _SC_EXPR_NEST_MAX      |
| {IOV_MAX}                                                                   | _SC_IOV_MAX            |
| {LINE_MAX}                                                                  | _SC_LINE_MAX           |
| {LOGIN_NAME_MAX}                                                            | _SC_LOGIN_NAME_MAX     |
| {NGROUPS_MAX}                                                               | _SC_NGROUPS_MAX        |
| Maximum size of <i>getgrgid_r()</i> and<br><i>getgrnam_r()</i> data buffers | _SC_GETGR_R_SIZE_MAX   |
| Maximum size of <i>getpwuid_r()</i> and<br><i>getpwnam_r()</i> data buffers | _SC_GETPW_R_SIZE_MAX   |
| {MQ_OPEN_MAX}                                                               | _SC_MQ_OPEN_MAX        |
| {MQ_PRIO_MAX}                                                               | _SC_MQ_PRIO_MAX        |
| {OPEN_MAX}                                                                  | _SC_OPEN_MAX           |
| _POSIX_ADVISORY_INFO                                                        | _SC_ADVISORY_INFO      |
| _POSIX_BARRIERS                                                             | _SC_BARRIERS           |
| _POSIX_ASYNCHRONOUS_IO                                                      | _SC_ASYNCHRONOUS_IO    |

|           | Variable                     | Value of Name             |
|-----------|------------------------------|---------------------------|
| 45549     |                              |                           |
| 45550     |                              |                           |
| 45551     | _POSIX_BASE                  | _SC_BASE                  |
| 45552     | _POSIX_C_LANG_SUPPORT        | _SC_C_LANG_SUPPORT        |
| 45553     | _POSIX_C_LANG_SUPPORT_R      | _SC_C_LANG_SUPPORT_R      |
| 45554 CS  | _POSIX_CLOCK_SELECTION       | _SC_CLOCK_SELECTION       |
| 45555 CPT | _POSIX_CPUTIME               | _SC_CPUTIME               |
| 45556     | _POSIX_DEVICE_IO             | _SC_DEVICE_IO             |
| 45557     | _POSIX_DEVICE_SPECIFIC       | _SC_DEVICE_SPECIFIC       |
| 45558     | _POSIX_DEVICE_SPECIFIC_R     | _SC_DEVICE_SPECIFIC_R     |
| 45559     | _POSIX_FD_MGMT               | _SC_FD_MGMT               |
| 45560     | _POSIX_FIFO                  | _SC_FIFO                  |
| 45561     | _POSIX_FILE_ATTRIBUTES       | _SC_FILE_ATTRIBUTES       |
| 45562     | _POSIX_FILE_LOCKING          | _SC_FILE_LOCKING          |
| 45563     | _POSIX_FILE_SYSTEM           | _SC_FILE_SYSTEM           |
| 45564 FSC | _POSIX_FSYNC                 | _SC_FSYNC                 |
| 45565     | _POSIX_JOB_CONTROL           | _SC_JOB_CONTROL           |
| 45566 MF  | _POSIX_MAPPED_FILES          | _SC_MAPPED_FILES          |
| 45567 ML  | _POSIX_MEMLOCK               | _SC_MEMLOCK               |
| 45568 MLR | _POSIX_MEMLOCK_RANGE         | _SC_MEMLOCK_RANGE         |
| 45569 MPR | _POSIX_MEMORY_PROTECTION     | _SC_MEMORY_PROTECTION     |
| 45570 MSG | _POSIX_MESSAGE_PASSING       | _SC_MESSAGE_PASSING       |
| 45571 MON | _POSIX_MONOTONIC_CLOCK       | _SC_MONOTONIC_CLOCK       |
| 45572     | _POSIX_MULTIPLE_PROCESS      | _SC_MULTIPLE_PROCESS      |
| 45573     | _POSIX_NETWORKING            | _SC_NETWORKING            |
| 45574     | _POSIX_PIPE                  | _SC_PIPE                  |
| 45575 PIO | _POSIX_PRIORITIZED_IO        | _SC_PRIORITIZED_IO        |
| 45576 PS  | _POSIX_PRIORITY_SCHEDULING   | _SC_PRIORITY_SCHEDULING   |
| 45577 THR | _POSIX_READER_WRITER_LOCKS   | _SC_READER_WRITER_LOCKS   |
| 45578 RTS | _POSIX_REALTIME_SIGNALS      | _SC_REALTIME_SIGNALS      |
| 45579     | _POSIX_REGEX                 | _SC_REGEX                 |
| 45580     | _POSIX_SAVED_IDS             | _SC_SAVED_IDS             |
| 45581 SEM | _POSIX_SEMAPHORES            | _SC_SEMAPHORES            |
| 45582 SHM | _POSIX_SHARED_MEMORY_OBJECTS | _SC_SHARED_MEMORY_OBJECTS |
| 45583     | _POSIX_SHELL                 | _SC_SHELL                 |
| 45584     | _POSIX_SIGNALS               | _SC_SIGNALS               |
| 45585     | _POSIX_SINGLE_PROCESS        | _SC_SINGLE_PROCESS        |
| 45586 SPN | _POSIX_SPAWN                 | _SC_SPAWN                 |
| 45587 SPI | _POSIX_SPIN_LOCKS            | _SC_SPIN_LOCKS            |
| 45588 SS  | _POSIX_SPORADIC_SERVER       | _SC_SPORADIC_SERVER       |
| 45589 SIO | _POSIX_SYNCHRONIZED_IO       | _SC_SYNCHRONIZED_IO       |
| 45590     | _POSIX_SYSTEM_DATABASE       | _SC_SYSTEM_DATABASE       |
| 45591     | _POSIX_SYSTEM_DATABASE_R     | _SC_SYSTEM_DATABASE_R     |

|           | Variable                          | Value of Name                  |
|-----------|-----------------------------------|--------------------------------|
| 45592     |                                   |                                |
| 45593     |                                   |                                |
| 45594 TSA | _POSIX_THREAD_ATTR_STACKADDR      | _SC_THREAD_ATTR_STACKADDR      |
| 45595 TSS | _POSIX_THREAD_ATTR_STACKSIZE      | _SC_THREAD_ATTR_STACKSIZE      |
| 45596 TCT | _POSIX_THREAD_CPUTIME             | _SC_THREAD_CPUTIME             |
| 45597 TPI | _POSIX_THREAD_PRIO_INHERIT        | _SC_THREAD_PRIO_INHERIT        |
| 45598 TPP | _POSIX_THREAD_PRIO_PROTECT        | _SC_THREAD_PRIO_PROTECT        |
| 45599 TPS | _POSIX_THREAD_PRIORITY_SCHEDULING | _SC_THREAD_PRIORITY_SCHEDULING |
| 45600 TSH | _POSIX_THREAD_PROCESS_SHARED      | _SC_THREAD_PROCESS_SHARED      |
| 45601 TSF | _POSIX_THREAD_SAFE_FUNCTIONS      | _SC_THREAD_SAFE_FUNCTIONS      |
| 45602 TSP | _POSIX_THREAD_SPORADIC_SERVER     | _SC_THREAD_SPORADIC_SERVER     |
| 45603 THR | _POSIX_THREADS                    | _SC_THREADS                    |
| 45604 TMO | _POSIX_TIMEOUTS                   | _SC_TIMEOUTS                   |
| 45605 TMR | _POSIX_TIMERS                     | _SC_TIMERS                     |
| 45606 TRC | _POSIX_TRACE                      | _SC_TRACE                      |
| 45607 TEF | _POSIX_TRACE_EVENT_FILTER         | _SC_TRACE_EVENT_FILTER         |
| 45608 TRI | _POSIX_TRACE_INHERIT              | _SC_TRACE_INHERIT              |
| 45609 TRL | _POSIX_TRACE_LOG                  | _SC_TRACE_LOG                  |
| 45610 TYM | _POSIX_TYPED_MEMORY_OBJECTS       | _SC_TYPED_MEMORY_OBJECTS       |
| 45611     | _POSIX_USER_GROUPS                | _SC_USER_GROUPS                |
| 45612     | _POSIX_USER_GROUPS_R              | _SC_USER_GROUPS_R              |
| 45613     | _POSIX_VERSION                    | _SC_VERSION                    |
| 45614     | _POSIX_V6_ILP32_OFF32             | _SC_V6_ILP32_OFF32             |
| 45615     | _POSIX_V6_ILP32_OFFBIG            | _SC_V6_ILP32_OFFBIG            |
| 45616     | _POSIX_V6_LP64_OFF64              | _SC_V6_LP64_OFF64              |
| 45617     | _POSIX_V6_LPBIG_OFFBIG            | _SC_V6_LPBIG_OFFBIG            |
| 45618     | _POSIX2_C_BIND                    | _SC_2_C_BIND                   |
| 45619     | _POSIX2_C_DEV                     | _SC_2_C_DEV                    |
| 45620     | _POSIX2_C_VERSION                 | _SC_2_C_VERSION                |
| 45621     | _POSIX2_CHAR_TERM                 | _SC_2_CHAR_TERM                |
| 45622     | _POSIX2_FORT_DEV                  | _SC_2_FORT_DEV                 |
| 45623     | _POSIX2_FORT_RUN                  | _SC_2_FORT_RUN                 |
| 45624     | _POSIX2_LOCALEDEF                 | _SC_2_LOCALEDEF                |
| 45625 BE  | _POSIX2_PBS                       | _SC_2_PBS                      |
| 45626     | _POSIX2_PBS_ACCOUNTING            | _SC_2_PBS_ACCOUNTING           |
| 45627     | _POSIX2_PBS_LOCATE                | _SC_2_PBS_LOCATE               |
| 45628     | _POSIX2_PBS_MESSAGE               | _SC_2_PBS_MESSAGE              |
| 45629     | _POSIX2_PBS_TRACK                 | _SC_2_PBS_TRACK                |
| 45630     | _POSIX2_SW_DEV                    | _SC_2_SW_DEV                   |
| 45631     | _POSIX2_UPE                       | _SC_2_UPE                      |
| 45632     | _POSIX2_VERSION                   | _SC_2_VERSION                  |
| 45633     | _REGEX_VERSION                    | _SC_REGEX_VERSION              |
| 45634 XSI | {PAGE_SIZE}                       | _SC_PAGE_SIZE                  |
| 45635     | {PAGESIZE}                        | _SC_PAGESIZE                   |

45636

45637

45638 THR

45639

45640

45641

45642

45643 RTS

45644 SEM

45645

45646 RTS

45647

45648

45649 TMR

45650

45651

45652 XSI

45653

45654

45655

45656

45657

45658

45659

45660

45661

45662

45663

45664

|  | Variable                        | Value of Name                    |
|--|---------------------------------|----------------------------------|
|  | {PTHREAD_DESTRUCTOR_ITERATIONS} | _SC_THREAD_DESTRUCTOR_ITERATIONS |
|  | {PTHREAD_KEYS_MAX}              | _SC_THREAD_KEYS_MAX              |
|  | {PTHREAD_STACK_MIN}             | _SC_THREAD_STACK_MIN             |
|  | {PTHREAD_THREADS_MAX}           | _SC_THREAD_THREADS_MAX           |
|  | {RE_DUP_MAX}                    | _SC_RE_DUP_MAX                   |
|  | {RTSIG_MAX}                     | _SC_RTSIG_MAX                    |
|  | {SEM_NSEMS_MAX}                 | _SC_SEM_NSEMS_MAX                |
|  | {SEM_VALUE_MAX}                 | _SC_SEM_VALUE_MAX                |
|  | {SIGQUEUE_MAX}                  | _SC_SIGQUEUE_MAX                 |
|  | {STREAM_MAX}                    | _SC_STREAM_MAX                   |
|  | {SYMLOOP_MAX}                   | _SC_SYMLOOP_MAX                  |
|  | {TIMER_MAX}                     | _SC_TIMER_MAX                    |
|  | {TTY_NAME_MAX}                  | _SC_TTY_NAME_MAX                 |
|  | {TZNAME_MAX}                    | _SC_TZNAME_MAX                   |
|  | _XBS5_ILP32_OFF32 (LEGACY)      | _SC_XBS5_ILP32_OFF32 (LEGACY)    |
|  | _XBS5_ILP32_OFFBIG (LEGACY)     | _SC_XBS5_ILP32_OFFBIG (LEGACY)   |
|  | _XBS5_LP64_OFF64 (LEGACY)       | _SC_XBS5_LP64_OFF64 (LEGACY)     |
|  | _XBS5_LPBIG_OFFBIG (LEGACY)     | _SC_XBS5_LPBIG_OFFBIG (LEGACY)   |
|  | _XOPEN_CRYPT                    | _SC_XOPEN_CRYPT                  |
|  | _XOPEN_ENH_I18N                 | _SC_XOPEN_ENH_I18N               |
|  | _XOPEN_LEGACY                   | _SC_XOPEN_LEGACY                 |
|  | _XOPEN_REALTIME                 | _SC_XOPEN_REALTIME               |
|  | _XOPEN_REALTIME_THREADS         | _SC_XOPEN_REALTIME_THREADS       |
|  | _XOPEN_SHM                      | _SC_XOPEN_SHM                    |
|  | _XOPEN_UNIX                     | _SC_XOPEN_UNIX                   |
|  | _XOPEN_VERSION                  | _SC_XOPEN_VERSION                |
|  | _XOPEN_XCU_VERSION              | _SC_XOPEN_XCU_VERSION            |

45665 **RETURN VALUE**

45666

45667

45668

If *name* is an invalid value, *sysconf()* shall return  $-1$  and set *errno* to indicate the error. If the variable corresponding to *name* has no limit, *sysconf()* shall return  $-1$  without changing the value of *errno*. Note that indefinite limits do not imply infinite limits; see <**limits.h**>.

45669

45670

45671

45672

Otherwise, *sysconf()* shall return the current variable value on the system. The value returned shall not be more restrictive than the corresponding value described to the application when it was compiled with the implementation's <**limits.h**> or <**unistd.h**>. The shall does not change during the lifetime of the calling process.

45673 **ERRORS**

45674

45675

The *sysconf()* function shall fail if:

[EINVAL]           The value of the *name* argument is invalid.

45676 **EXAMPLES**

45677

None.

45678 **APPLICATION USAGE**

45679

45680

45681

As  $-1$  is a permissible return value in a successful situation, an application wishing to check for error situations should set *errno* to 0, then call *sysconf()*, and, if it returns  $-1$ , check to see if *errno* is non-zero.

45682

45683

45684

If the value of *sysconf(\_SC\_2\_VERSION)* is not equal to the value of the *\_POSIX2\_VERSION* symbolic constant, the utilities available via *system()* or *popen()* might not behave as described in the Shell and Utilities volume of IEEE Std. 1003.1-200x. This would mean that the application is

45685 not running in an environment that conforms to the Shell and Utilities volume of  
45686 IEEE Std. 1003.1-200x. Some applications might be able to deal with this, others might not.  
45687 However, the functions defined in this volume of IEEE Std. 1003.1-200x continue to operate as  
45688 specified, even if: `sysconf(_SC_2_VERSION)` reports that the utilities no longer perform as  
45689 specified.

#### 45690 RATIONALE

45691 This functionality was added in response to requirements of application developers and of  
45692 system vendors who deal with many international system configurations. It is closely related to  
45693 `pathconf()` and `fpathconf()`.

45694 Although a portable application can run on all systems by never demanding more resources  
45695 than the minimum values published in this volume of IEEE Std. 1003.1-200x, it is useful for that  
45696 application to be able to use the actual value for the quantity of a resource available on any  
45697 given system. To do this, the application makes use of the value of a symbolic constant in  
45698 `<limits.h>` or `<unistd.h>`.

45699 However, once compiled, the application must still be able to cope if the amount of resource  
45700 available is increased. To that end, an application may need a means of determining the quantity  
45701 of a resource, or the presence of an option, at execution time.

45702 Two examples are offered:

- 45703 1. Applications may wish to act differently on systems with or without job control.  
45704 Applications vendors who wish to distribute only a single binary package to all instances  
45705 of a computer architecture would be forced to assume job control is never available if it  
45706 were to rely solely on the `<unistd.h>` value published in this volume of  
45707 IEEE Std. 1003.1-200x.
- 45708 2. International applications vendors occasionally require knowledge of the number of clock  
45709 ticks per second. Without these facilities, they would be required to either distribute their  
45710 applications partially in source form or to have 50Hz and 60Hz versions for the various  
45711 countries in which they operate.

45712 It is the knowledge that many applications are actually distributed widely in executable form  
45713 that leads to this facility. If limited to the most restrictive values in the headers, such  
45714 applications would have to be prepared to accept the most limited environments offered by the  
45715 smallest microcomputers. Although this is entirely portable, there was a consensus that they  
45716 should be able to take advantage of the facilities offered by large systems, without the  
45717 restrictions associated with source and object distributions.

45718 During the discussions of this feature, it was pointed out that it is almost always possible for an  
45719 application to discern what a value might be at runtime by suitably testing the various functions  
45720 themselves. And, in any event, it could always be written to adequately deal with error returns  
45721 from the various functions. In the end, it was felt that this imposed an unreasonable level of  
45722 complication and sophistication on the application writer.

45723 This runtime facility is not meant to provide ever-changing values that applications have to  
45724 check multiple times. The values are seen as changing no more frequently than once per system  
45725 initialization, such as by a system administrator or operator with an automatic configuration  
45726 program. This volume of IEEE Std. 1003.1-200x specifies that they shall not change within the  
45727 lifetime of the process.

45728 Some values apply to the system overall and others vary at the file system or directory level. The  
45729 latter are described in `pathconf()`.

45730 Note that all values returned must be expressible as integers. String values were considered, but  
45731 the additional flexibility of this approach was rejected due to its added complexity of

45732 implementation and use.

45733 Some values, such as {PATH\_MAX}, are sometimes so large that they must not be used to, say,  
45734 allocate arrays. The *sysconf()* function returns a negative value to show that this symbolic  
45735 constant is not even defined in this case.

45736 Similar to *pathconf()*, this permits the implementation not to have a limit. When one resource is  
45737 infinite, returning an error indicating that some other resource limit has been reached is  
45738 conforming behavior.

#### 45739 FUTURE DIRECTIONS

45740 None.

#### 45741 SEE ALSO

45742 *confstr()*, *pathconf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <limits.h>,  
45743 <unistd.h>, the Shell and Utilities volume of IEEE Std. 1003.1-200x, *getconf*

#### 45744 CHANGE HISTORY

45745 First released in Issue 3.

45746 Entry included for alignment with the POSIX.1-1988 standard.

#### 45747 Issue 4

45748 The type of the function return value is expanded to **long**.

45749 `_XOPEN_VERSION` is added to the table of configurable system limits; this should have been  
45750 included in Issue 3.

45751 The following variables are added to the table of configurable system limits in the  
45752 DESCRIPTION and marked as extensions:

45753 `_XOPEN_CRYPT`  
45754 `_XOPEN_ENH_I18N`  
45755 `_XOPEN_SHM`  
45756 `_XOPEN_UNIX`

45757 In the RETURN VALUE section the header <time.h> is given as an alternative to <limits.h> and  
45758 <unistd.h>.

45759 The second paragraph is added to the APPLICATION USAGE section.

45760 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 45761 • The variables {STREAM\_MAX} and {TZNAME\_MAX} are added to the table of variables in  
45762 the DESCRIPTION.

45763 The following change is incorporated for alignment with the ISO POSIX-2 standard:

- 45764 • The following variables are added to the table of configurable system limits in the  
45765 DESCRIPTION:

|       |                    |                   |                 |
|-------|--------------------|-------------------|-----------------|
| 45766 | {BC_BASE_MAX}      | _POSIX2_C_BIND    | _POSIX2_SW_DEV  |
| 45767 | {BC_DIM_MAX}       | _POSIX2_C_DEV     | _POSIX2_VERSION |
| 45768 | {BC_SCALE_MAX}     | _POSIX2_C_VERSION | {RE_DUP_MAX}    |
| 45769 | {BC_STRING_MAX}    | _POSIX2_CHAR_TERM |                 |
| 45770 | {COLL_WEIGHTS_MAX} | _POSIX2_FORT_DEV  |                 |
| 45771 | {EXPR_NEST_MAX}    | _POSIX2_FORT_RUN  |                 |
| 45772 | {LINE_MAX}         | _POSIX2_LOCALEDEF |                 |

**45773 Issue 4, Version 2**

45774 For X/OPEN UNIX conformance, the {ATEXIT\_MAX}, {IOV\_MAX}, {PAGESIZE}, {PAGE\_SIZE},  
 45775 and \_XOPEN\_UNIX variables are added to the list of configurable system values that can be  
 45776 determined by calling *sysconf()*.

**45777 Issue 5**

45778 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX  
 45779 Threads Extension.

45780 The \_XBS\_ variables and name values are added to the table of system variables in the  
 45781 DESCRIPTION. These are all marked EX.

**45782 Issue 6**

45783 The symbol CLK\_TCK is obsolescent and removed. It is replaced with the phrase “clock ticks  
 45784 per second”.

45785 The symbol {PASS\_MAX} is removed.

45786 The following changes were made to align with the IEEE P1003.1a draft standard:

45787 • Table entries added for the following variables: \_SC\_REGEX, \_SC\_SHELL,  
 45788 \_SC\_REGEX\_VERSION, \_SC\_SYMLINK\_MAX.

45789 The following *sysconf()* variables and their associated names are added for alignment with  
 45790 IEEE Std. 1003.1d-1999:

45791 \_POSIX\_ADVISORY\_INFO  
 45792 \_POSIX\_CPUTIME  
 45793 \_POSIX\_SPAWN  
 45794 \_POSIX\_SPORADIC\_SERVER  
 45795 \_POSIX\_THREAD\_CPUTIME  
 45796 \_POSIX\_THREAD\_SPORADIC\_SERVER  
 45797 \_POSIX\_TIMEOUTS

45798 The following changes are made to the DESCRIPTION for alignment with  
 45799 IEEE Std. 1003.1j-2000:

45800 • A statement expressing the dependency of support for some system variables on  
 45801 implementation options is added.

45802 • The following system variables are added:

45803 \_POSIX\_BARRIERS  
 45804 \_POSIX\_CLOCK\_SELECTION  
 45805 \_POSIX\_MONOTONIC\_CLOCK  
 45806 \_POSIX\_READER\_WRITER\_LOCKS  
 45807 \_POSIX\_SPIN\_LOCKS  
 45808 \_POSIX\_TYPED\_MEMORY\_OBJECTS

45809 The following system variables are added for alignment with IEEE Std. 1003.2d-1994:

45810 \_POSIX2\_PBS  
 45811 \_POSIX2\_PBS\_ACCOUNTING  
 45812 \_POSIX2\_PBS\_LOCATE  
 45813 \_POSIX2\_PBS\_MESSAGE  
 45814 \_POSIX2\_PBS\_TRACK

45815 The following *sysconf()* variables and their associated names are added for alignment with  
 45816 IEEE Std. 1003.1q-2000:

45817            \_POSIX\_TRACE  
45818            \_POSIX\_TRACE\_EVENT\_FILTER  
45819            \_POSIX\_TRACE\_INHERIT  
45820            \_POXIC\_TRACE\_LOG

45821            The macros associated with the *c89* programming models are marked LEGACY, and new  
45822            equivalent macros associated with *c99* are introduced.

45823 **NAME**

45824        syslog — log a message

45825 **SYNOPSIS**

45826 XSI     #include <syslog.h>

45827        void syslog(int *priority*, const char \**message*, ... /\* *argument* \*/);

45828

45829 **DESCRIPTION**

45830        Refer to *closelog*().

45831 **NAME**

45832 system — issue a command

45833 **SYNOPSIS**

45834 #include &lt;stdlib.h&gt;

45835 int system(const char \**command*);45836 **DESCRIPTION**

45837 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 45838 conflict between the requirements described here and the ISO C standard is unintentional. This  
 45839 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

45840 The *system()* function passes the string pointed to by *command* to the host environment to be  
 45841 executed by a command processor in an implementation-defined manner. The environment of  
 45842 the executed command shall be as if a child process were created using the *fork()* function, and  
 45843 the child process invoked a command interpreter using the *execl()* function.

45844 CX If the implementation supports the Shell and Utilities volume of IEEE Std. 1003.1-200x  
 45845 commands, the environment of the executed command shall be as if a child process were created  
 45846 using *fork()*, and the child process invoked the *sh* utility using *execl()* as follows:

```
45847 execl(<shell path>, "sh", "-c", command, (char *)0);
```

45848 where <shell path> is an unspecified path name for the *sh* utility.

45849 The *system()* function ignores the SIGINT and SIGQUIT signals, and blocks the SIGCHLD  
 45850 signal, while waiting for the command to terminate. If this might cause the application to miss a  
 45851 signal that would have killed it, then the application should examine the return value from  
 45852 *system()* and take whatever action is appropriate to the application if the command terminated  
 45853 due to receipt of a signal.

45854 The *system()* function shall not affect the termination status of any child of the calling processes  
 45855 other than the process or processes it itself creates.

45856 The *system()* function shall not return until the child process has terminated.

45857 **RETURN VALUE**

45858 If *command* is a null pointer, *system()* shall return non-zero to indicate that a command processor  
 45859 CX is available, or zero if none is available. If the implementation supports the utilities defined in  
 45860 the Shell and Utilities volume of IEEE Std. 1003.1-200x, *system()* shall always return non-zero  
 45861 when *command* is NULL.

45862 CX If *command* is not a null pointer, *system()* shall return the termination status of the command  
 45863 language interpreter in the format specified by *waitpid()*. If the implementation supports the  
 45864 utilities defined in the Shell and Utilities volume of IEEE Std. 1003.1-200x, the termination status  
 45865 shall be as defined for the *sh* utility; otherwise, the termination status is unspecified. If some  
 45866 error prevents the command language interpreter from executing after the child process is  
 45867 created, the return value from *system()* shall be as if the command language interpreter had  
 45868 terminated using *exit(127)* or *\_exit(127)*. If a child process cannot be created, or if the  
 45869 termination status for the command language interpreter cannot be obtained, *system()* shall  
 45870 return -1 and set *errno* to indicate the error.

45871 **ERRORS**

45872 CX The *system()* function may set *errno* values as described by *fork()*.

45873 In addition, *system()* may fail if:

45874 CX [ECHILD] The status of the child process created by *system()* is no longer available.

45875 **EXAMPLES**

45876 None.

45877 **APPLICATION USAGE**

45878 If the return value of *system()* is not `-1`, its value can be decoded through the use of the macros  
45879 described in `<sys/wait.h>`. For convenience, these macros are also provided in `<stdlib.h>`.

45880 To determine whether or not the environment specified in the Shell and Utilities volume of  
45881 IEEE Std. 1003.1-200x is present, use *sysconf(SC\_2\_VERSION)*.

45882 Note that, while *system()* must ignore SIGINT and SIGQUIT and block SIGCHLD while waiting  
45883 for the child to terminate, the handling of signals in the executed command is as specified by  
45884 *fork()* and *exec*. For example, if SIGINT is being caught or is set to SIG\_DFL when *system()* is  
45885 called, then the child is started with SIGINT handling set to SIG\_DFL.

45886 Ignoring SIGINT and SIGQUIT in the parent process prevents coordination problems (two  
45887 processes reading from the same terminal, for example) when the executed command ignores or  
45888 catches one of the signals. It is also usually the correct action when the user has given a  
45889 command to the application to be executed synchronously (as in the `'!'` command in many  
45890 interactive applications). In either case, the signal should be delivered only to the child process,  
45891 not to the application itself. There is one situation where ignoring the signals might have less  
45892 than the desired effect. This is when the application uses *system()* to perform some task invisible  
45893 to the user. If the user typed the interrupt character ("`^C`", for example) while *system()* is being  
45894 used in this way, one would expect the application to be killed, but only the executed command  
45895 is killed. Applications that use *system()* in this way should carefully check the return status from  
45896 *system()* to see if the executed command was successful, and should take appropriate action  
45897 when the command fails.

45898 Blocking SIGCHLD while waiting for the child to terminate prevents the application from  
45899 catching the signal and obtaining status from *system()*'s child process before *system()* can get the  
45900 status itself.

45901 The context in which the utility is ultimately executed may differ from that in which *system()*  
45902 was called. For example, file descriptors that have the FD\_CLOEXEC flag set are closed, and the  
45903 process ID and parent process ID are different. Also, if the executed utility changes its  
45904 environment variables or its current working directory, that change is not reflected in the caller's  
45905 context.

45906 There is no defined way for an application to find the specific path for the shell. However,  
45907 *confstr()* can provide a value for *PATH* that is guaranteed to find the *sh* utility.

45908 **RATIONALE**

45909 The *system()* function should not be used by programs that have set user (or group) ID  
45910 privileges. The *fork()* and *exec* family of functions (except *execlp()* and *execvp()*), should be used  
45911 instead. This prevents any unforeseen manipulation of the environment of the user that could  
45912 cause execution of commands not anticipated by the calling program.

45913 There are three levels of specification for the *system()* function. The ISO C standard gives the  
45914 most basic. It requires that the function exists, and defines a way for an application to query  
45915 whether a command language interpreter exists. It says nothing about the command language or  
45916 the environment in which the command is interpreted.

45917 IEEE Std. 1003.1-200x places additional restrictions on *system()*. It requires that if there is a  
45918 command language interpreter, the environment must be as specified by *fork()* and *exec*. This  
45919 ensures, for example, that close-on-exec works, that file locks are not inherited, and that the  
45920 process ID is different. It also specifies the return value from *system()* when the command line  
45921 can be run, thus giving the application some information about the command's completion

45922 status IEEE Std. 1003.1-200x in its base definition still says nothing about the interpretation of  
45923 the command.

45924 Finally, IEEE Std. 1003.1-200x requires the command to be interpreted as in the shell command  
45925 language defined in the Shell and Utilities volume of IEEE Std. 1003.1-200x.

45926 Note that, *system*(NULL) is required to return non-zero, indicating that there is a command  
45927 language interpreter. At first glance, this would seem to conflict with the ISO C standard which  
45928 allows *system*(NULL) to return zero. There is no conflict, however. A system must have a  
45929 command language interpreter, and is non-conforming if none is present. It is therefore  
45930 permissible for the *system*() function on such a system to implement the behavior specified by  
45931 the ISO C standard as long as it is understood that the implementation does not conform to  
45932 IEEE Std. 1003.1-200x if *system*(NULL) returns zero.

45933 It was explicitly decided that when *command* is NULL, *system*() should not be required to check  
45934 to make sure that the command language interpreter actually exists with the correct mode, that  
45935 there are enough processes to execute it, and so on. The call *system*(NULL) could, theoretically,  
45936 check for such problems as too many existing child processes, and return zero. However, it  
45937 would be inappropriate to return zero due to such a (presumably) transient condition. If some  
45938 condition exists that is not under the control of this application and that would cause any  
45939 *system*() call to fail, that system has been rendered non-conforming.

45940 Early drafts required, or allowed, *system*() to return with *errno* set to [EINTR] if it was  
45941 interrupted with a signal. This error return was removed, and a requirement that *system*() not  
45942 return until the child has terminated was added. This means that if a *waitpid*() call in *system*()  
45943 exits with *errno* set to [EINTR], *system*() must re-issue the *waitpid*(). This change was made for  
45944 two reasons:

- 45945 1. There is no way for an application to clean up if *system*() returns [EINTR], short of calling  
45946 *wait*(), and that could have the undesirable effect of returning the status of children other  
45947 than the one started by *system*()).
- 45948 2. While it might require a change in some historical implementations, those  
45949 implementations already have to be changed because they use *wait*() instead of *waitpid*()).

45950 Note that if the application is catching SIGCHLD signals, it will receive such a signal before a  
45951 successful *system*() call returns.

45952 To conform to IEEE Std. 1003.1-200x, *system*() must use *waitpid*(), or some similar function,  
45953 instead of *wait*()).

45954 The following code sample illustrates how *system*() might be implemented on an  
45955 implementation conforming to IEEE Std. 1003.1-200x.

```
45956 #include <signal.h>
45957 int system(const char *cmd)
45958 {
45959 int stat;
45960 pid_t pid;
45961 struct sigaction sa, savintr, savequit;
45962 sigset_t saveblock;
45963 if (cmd == NULL)
45964 return(1);
45965 sa.sa_handler = SIG_IGN;
45966 sigemptyset(&sa.sa_mask);
45967 sa.sa_flags = 0;
45968 sigemptyset(&saveintr.sa_mask);
```

```

45969 sigemptyset(&savequit.sa_mask);
45970 sigaction(SIGINT, &sa, &saveintr);
45971 sigaction(SIGQUIT, &sa, &savequit);
45972 sigaddset(&sa.sa_mask, SIGCHLD);
45973 sigprocmask(SIG_BLOCK, &sa.sa_mask, &saveblock);
45974 if ((pid = fork()) == 0) {
45975 sigaction(SIGINT, &saveintr, (struct sigaction *)0);
45976 sigaction(SIGQUIT, &savequit, (struct sigaction *)0);
45977 sigprocmask(SIG_SETMASK, &saveblock, (sigset_t *)0);
45978 execl("/bin/sh", "sh", "-c", cmd, (char *)0);
45979 _exit(127);
45980 }
45981 if (pid == -1) {
45982 stat = -1; /* errno comes from fork() */
45983 } else {
45984 while (waitpid(pid, &stat, 0) == -1) {
45985 if (errno != EINTR){
45986 stat = -1;
45987 break;
45988 }
45989 }
45990 }
45991 sigaction(SIGINT, &saveintr, (struct sigaction *)0);
45992 sigaction(SIGQUIT, &savequit, (struct sigaction *)0);
45993 sigprocmask(SIG_SETMASK, &saveblock, (sigset_t *)0);
45994 return(stat);
45995 }

```

45996 Note that, while a particular implementation of *system()* (such as the one above) can assume a  
45997 particular path for the shell, such a path is not necessarily valid on another system. The above  
45998 example is not portable, and is not intended to be.

45999 One reviewer suggested that an implementation of *system()* might want to use an environment  
46000 variable such as *SHELL* to determine which command interpreter to use. The supposed  
46001 implementation would use the default command interpreter if the one specified by the  
46002 environment variable was not available. This would allow a user, when using an application  
46003 that prompts for command lines to be processed using *system()*, to specify a different command  
46004 interpreter. Such an implementation is discouraged. If the alternate command interpreter did not  
46005 follow the command line syntax specified in the Shell and Utilities volume of  
46006 IEEE Std. 1003.1-200x, then changing *SHELL* would render *system()* non-conforming. This would  
46007 affect applications that expected the specified behavior from *system()*, and since the Shell and  
46008 Utilities volume of IEEE Std. 1003.1-200x does not mention that *SHELL* affects *system()*, the  
46009 application would not know that it needed to unset *SHELL*.

#### 46010 FUTURE DIRECTIONS

46011 None.

#### 46012 SEE ALSO

46013 *exec*, *pipe()*, *waitpid()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**limits.h**>,  
46014 <**signal.h**>, <**stdlib.h**>, <**sys/wait.h**>, the Shell and Utilities volume of IEEE Std. 1003.1-200x

46015 **CHANGE HISTORY**

46016 First released in Issue 1. Derived from Issue 1 of the SVID.

46017 **Issue 4**

46018 Extensions beyond the ISO C standard are now marked.

46019 The following changes are incorporated for alignment with the ISO POSIX-2 standard:

- 46020
  - The function is no longer marked as an extension.
- 46021
  - The name of the argument is changed from *string* to *command*, and its type is changed from
- 46022
  - char\*** to **const char\***.
- 46023
  - The DESCRIPTION and RETURN VALUE sections are completely replaced to bring them in
- 46024
  - line with the ISO POSIX-2 standard. They still describe essentially the same functionality,
- 46025
  - albeit that the definition is more complete.
- 46026
  - The ERRORS section is changed to indicate that *system()* may return error values described
- 46027
  - for *fork()*.
- 46028
  - The APPLICATION USAGE section is added.

46029 The following changes were made to align with the IEEE P1003.1a draft standard:

- 46030
  - The DESCRIPTION is adjusted to reflect the behavior on systems that do not support the
- 46031
  - Shell option.

46032 **NAME**

46033 tan, tanf, tanl — tangent function

46034 **SYNOPSIS**

46035 #include &lt;math.h&gt;

46036 double tan(double x);

46037 float tanf(float x);

46038 long double tanl(long double x);

46039 **DESCRIPTION**

46040 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 46041 conflict between the requirements described here and the ISO C standard is unintentional. This  
 46042 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

46043 These functions shall compute the tangent of its argument *x*, measured in radians.

46044 An application wishing to check for error situations should set *errno* to 0 before calling *tan()*. If  
 46045 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

46046 The *tan()* function may lose accuracy when its argument is far from 0.0.46047 **RETURN VALUE**46048 Upon successful completion, these functions shall return the tangent of *x*.46049 XSI If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

46050 If *x* is  $\pm\text{Inf}$ , either 0.0 shall be returned and *errno* set to [EDOM], or NaN shall be returned and  
 46051 *errno* may be set to [EDOM].

46052 If the correct value would cause overflow,  $\pm\text{HUGE\_VAL}$  shall be returned and *errno* shall be set  
 46053 to [ERANGE].

46054 If the correct value would cause underflow, 0.0 shall be returned and *errno* may be set to  
 46055 [ERANGE].

46056 **ERRORS**

46057 These functions shall fail if:

46058 [ERANGE] The value to be returned would cause overflow.

46059 These functions may fail if:

46060 XSI [EDOM] The value *x* is NaN or  $\pm\text{Inf}$ .

46061 [ERANGE] The value to be returned would cause underflow.

46062 XSI No other errors shall occur.

46063 **EXAMPLES**46064 **Taking the Tangent of a 45-Degree Angle**

46065 #include &lt;math.h&gt;

46066 ...

46067 double radians = 45.0 \* M\_PI / 180;

46068 double result;

46069 ...

46070 result = tan (radians);

46071 **APPLICATION USAGE**

46072 None.

46073 **RATIONALE**

46074 None.

46075 **FUTURE DIRECTIONS**

46076 None.

46077 **SEE ALSO**46078 *atan()*, *isnan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>46079 **CHANGE HISTORY**

46080 First released in Issue 1. Derived from Issue 1 of the SVID.

46081 **Issue 4**46082 References to *matherr()* are removed.46083 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
46084 ISO C standard and to rationalize error handling in the mathematics functions.

46085 The return value specified for [EDOM] is marked as an extension.

46086 **Issue 5**46087 The last two paragraphs of the DESCRIPTION were included as APPLICATION USAGE notes  
46088 in previous issues.46089 **Issue 6**46090 The *tanf()* and *tanl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

46091 **NAME**

46092 tanh, tanhf, tanhl — hyperbolic tangent function

46093 **SYNOPSIS**

46094 #include &lt;math.h&gt;

46095 double tanh(double x);

46096 float tanhf(float x);

46097 long double tanhl(long double x);

46098 **DESCRIPTION**

46099 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
46100 conflict between the requirements described here and the ISO C standard is unintentional. This  
46101 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

46102 These functions shall compute the hyperbolic tangent of  $x$ .

46103 An application wishing to check for error situations should set *errno* to 0 before calling *tanh()*. If  
46104 *errno* is non-zero on return, or the return value is NaN, an error has occurred.

46105 **RETURN VALUE**46106 Upon successful completion, these functions shall return the hyperbolic tangent of  $x$ .46107 XSI If  $x$  is NaN, NaN shall be returned and *errno* may be set to [EDOM].

46108 If the correct value would cause underflow, 0.0 shall be returned and *errno* may be set to  
46109 [ERANGE].

46110 **ERRORS**

46111 These functions may fail if:

46112 XSI [EDOM] The value of  $x$  is NaN.

46113 [ERANGE] The correct result would cause underflow.

46114 XSI No other errors shall occur.

46115 **EXAMPLES**

46116 None.

46117 **APPLICATION USAGE**

46118 None.

46119 **RATIONALE**

46120 None.

46121 **FUTURE DIRECTIONS**

46122 None.

46123 **SEE ALSO**46124 *atanh()*, *isnan()*, *tan()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>46125 **CHANGE HISTORY**

46126 First released in Issue 1. Derived from Issue 1 of the SVID.

46127 **Issue 4**46128 References to *matherr()* are removed.

46129 The RETURN VALUE and ERRORS sections are substantially rewritten for alignment with the  
46130 ISO C standard and to rationalize error handling in the mathematics functions.

46131 The return value specified for [EDOM] is marked as an extension.

46132 **Issue 5**

46133 The DESCRIPTION is updated to indicate how an application should check for an error. This  
46134 text was previously published in the APPLICATION USAGE section.

46135 **Issue 6**

46136 The *tanhf()* and *tanhl()* functions are added for alignment with the ISO/IEC 9899:1999 standard.

46137 **NAME**

46138 tcdrain — wait for transmission of output

46139 **SYNOPSIS**

46140 #include &lt;termios.h&gt;

46141 int tcdrain(int *fildev*);46142 **DESCRIPTION**46143 The *tcdrain()* function shall wait until all output written to the object referred to by *fildev* is  
46144 transmitted. The *fildev* argument is an open file descriptor associated with a terminal.46145 Any attempts to use *tcdrain()* from a process which is a member of a background process group  
46146 on a *fildev* associated with its controlling terminal, shall cause the process group to be sent a  
46147 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is  
46148 allowed to perform the operation, and no signal is sent.46149 **RETURN VALUE**46150 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46151 indicate the error.46152 **ERRORS**46153 The *tcdrain()* function shall fail if:46154 [EBADF] The *fildev* argument is not a valid file descriptor. |46155 [EINTR] A signal interrupted *tcdrain()*. |46156 [ENOTTY] The file associated with *fildev* is not a terminal. |46157 The *tcdrain()* function may fail if:46158 [EIO] The process group of the writing process is orphaned, and the writing process |  
46159 is not ignoring or blocking SIGTTOU. |46160 **EXAMPLES**

46161 None.

46162 **APPLICATION USAGE**

46163 None.

46164 **RATIONALE**

46165 None.

46166 **FUTURE DIRECTIONS**

46167 None.

46168 **SEE ALSO**46169 *tcfldsh()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <termios.h>, <unistd.h>, the |  
46170 Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 11, General Terminal Interface |46171 **CHANGE HISTORY**

46172 First released in Issue 3.

46173 Entry included for alignment with the POSIX.1-1988 standard.

46174 **Issue 4**

46175 The [EIO] error is added to the ERRORS section.

46176 The FUTURE DIRECTIONS section is added.

46177 The following change is incorporated for alignment with the FIPS requirements:

- 46178           • The words “If \_POSIX\_JOB\_CONTROL is defined” are removed from the start of the second  
46179           paragraph in the DESCRIPTION. This is because job control is defined as mandatory for  
46180           Issue 4 conforming implementations.

46181 **Issue 6**

46182           The following new requirements on POSIX implementations derive from alignment with the  
46183           Single UNIX Specification:

- 46184           • In the DESCRIPTION, the final paragraph is no longer conditional on  
46185           \_POSIX\_JOB\_CONTROL. This is a FIPS requirement.
- 46186           • The [EIO] error is added.

46187 **NAME**

46188 tcflow — suspend or restart the transmission or reception of data

46189 **SYNOPSIS**

46190 #include &lt;termios.h&gt;

46191 int tcflow(int *fildev*, int *action*);46192 **DESCRIPTION**46193 The *tcflow()* function shall suspend or restart transmission or reception of data on the object  
46194 referred to by *fildev*, depending on the value of *action*. The *fildev* argument is an open file  
46195 descriptor associated with a terminal.

- 46196 • If *action* is TCOFF, output shall be suspended.
- 46197 • If *action* is TCOON, suspended output shall be restarted.
- 46198 • If *action* is TCIOFF, the system shall transmit a STOP character, which is intended to cause  
46199 the terminal device to stop transmitting data to the system.
- 46200 • If *action* is TCION, the system shall transmit a START character, which is intended to cause  
46201 the terminal device to start transmitting data to the system.

46202 The default on the opening of a terminal file is that neither its input nor its output are  
46203 suspended.46204 Attempts to use *tcflow()* from a process which is a member of a background process group on a  
46205 *fildev* associated with its controlling terminal, shall cause the process group to be sent a  
46206 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is  
46207 allowed to perform the operation, and no signal is sent.46208 **RETURN VALUE**46209 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46210 indicate the error.46211 **ERRORS**46212 The *tcflow()* function shall fail if:

- 46213 [EBADF] The *fildev* argument is not a valid file descriptor.
- 46214 [EINVAL] The *action* argument is not a supported value.
- 46215 [ENOTTY] The file associated with *fildev* is not a terminal.

46216 The *tcflow()* function may fail if:

- 46217 [EIO] The process group of the writing process is orphaned, and the writing process  
46218 is not ignoring or blocking SIGTTOU.

46219 **EXAMPLES**

46220 None.

46221 **APPLICATION USAGE**

46222 None.

46223 **RATIONALE**

46224 None.

46225 **FUTURE DIRECTIONS**

46226 None.

46227 **SEE ALSO**

46228 *tcsendbreak()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <termios.h>, <unistd.h>, the  
46229 Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 11, General Terminal Interface

46230 **CHANGE HISTORY**

46231 First released in Issue 3.

46232 Entry included for alignment with the POSIX.1-1988 standard.

46233 **Issue 4**

46234 The descriptions of TCIOFF and TCION are reworded, indicating the intended consequences of  
46235 transmitting stop and start characters. Issue 3 implied that these consequences were guaranteed.

46236 The [EIO] error is added to the ERRORS section.

46237 The FUTURE DIRECTIONS section is added.

46238 The following change is incorporated for alignment with the FIPS requirements:

- 46239 • The words “If \_POSIX\_JOB\_CONTROL is defined” are removed from the start of the second  
46240 paragraph in the DESCRIPTION. This is because job control is defined as mandatory for  
46241 Issue 4 conforming implementations.

46242 **Issue 6**

46243 The following new requirements on POSIX implementations derive from alignment with the  
46244 Single UNIX Specification:

- 46245 • The [EIO] error is added.

46246 **NAME**

46247 tcflush — flush non-transmitted output data, non-read input data, or both

46248 **SYNOPSIS**

46249 #include &lt;termios.h&gt;

46250 int tcflush(int *fildev*, int *queue\_selector*);46251 **DESCRIPTION**

46252 Upon successful completion, *tcflush()* shall discard data written to the object referred to by *fildev*  
 46253 (an open file descriptor associated with a terminal) but not transmitted, or data received but not  
 46254 read, depending on the value of *queue\_selector*:

- 46255 • If *queue\_selector* is TCIFLUSH, it shall flush data received but not read.
- 46256 • If *queue\_selector* is TCOFLUSH, it shall flush data written but not transmitted.
- 46257 • If *queue\_selector* is TCIOFLUSH, it shall flush both data received but not read and data  
 46258 written but not transmitted.

46259 Attempts to use *tcflush()* from a process which is a member of a background process group on a  
 46260 *fildev* associated with its controlling terminal, shall cause the process group to be sent a  
 46261 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is  
 46262 allowed to perform the operation, and no signal is sent.

46263 **RETURN VALUE**

46264 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
 46265 indicate the error.

46266 **ERRORS**46267 The *tcflush()* function shall fail if:46268 [EBADF] The *fildev* argument is not a valid file descriptor. |46269 [EINVAL] The *queue\_selector* argument is not a supported value. |46270 [ENOTTY] The file associated with *fildev* is not a terminal. |46271 The *tcflush()* function may fail if:

46272 [EIO] The process group of the writing process is orphaned, and the writing process |  
 46273 is not ignoring or blocking SIGTTOU. |

46274 **EXAMPLES**

46275 None.

46276 **APPLICATION USAGE**

46277 None.

46278 **RATIONALE**

46279 None.

46280 **FUTURE DIRECTIONS**

46281 None.

46282 **SEE ALSO**

46283 *tcdrain()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <termios.h>, <unistd.h>, the |  
 46284 Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 11, General Terminal Interface |

46285 **CHANGE HISTORY**

46286 First released in Issue 3.

46287 Entry included for alignment with the POSIX.1-1988 standard.

46288 **Issue 4**46289 The DESCRIPTION is modified to indicate that the flush operation only results if the call to  
46290 *tcflush()* is successful.

46291 The [EIO] error is added to the ERRORS section.

46292 The FUTURE DIRECTIONS section is added.

46293 The following change is incorporated for alignment with the FIPS requirements:

- 46294
- The words “If \_POSIX\_JOB\_CONTROL is defined” are removed from the start of the second  
46295 paragraph in the DESCRIPTION. This is because job control is defined as mandatory for  
46296 Issue 4 conforming implementations.

46297 **Issue 6**46298 The Open Group corrigenda item U035/1 has been applied. In the ERRORS and APPLICATION  
46299 USAGE sections, references to *tcflow()* are replaced with *tcflush()*.46300 The following new requirements on POSIX implementations derive from alignment with the  
46301 Single UNIX Specification:

- 46302
- In the DESCRIPTION, the final paragraph is no longer conditional on  
46303 \_POSIX\_JOB\_CONTROL. This is a FIPS requirement.

- 46304
- The [EIO] error is added.

## 46305 NAME

46306 tcgetattr — get the parameters associated with the terminal

## 46307 SYNOPSIS

46308 #include <termios.h>

46309 int tcgetattr(int *fildev*, struct termios \**termios\_p*);

## 46310 DESCRIPTION

46311 The *tcgetattr()* function shall get the parameters associated with the terminal referred to by *fildev*  
46312 and store them in the **termios** structure referenced by *termios\_p*. The *fildev* argument is an open  
46313 file descriptor associated with a terminal.

46314 The *termios\_p* argument is a pointer to a **termios** structure.

46315 The *tcgetattr()* operation is allowed from any process.

46316 If the terminal device supports different input and output baud rates, the baud rates stored in  
46317 the **termios** structure returned by *tcgetattr()* shall reflect the actual baud rates, even if they are  
46318 equal. If differing baud rates are not supported, the rate returned as the output baud rate shall be  
46319 the actual baud rate. If the terminal device does not support split baud rates, the input baud rate  
46320 stored in the **termios** structure shall be the output rate (as one of the symbolic values).

## 46321 RETURN VALUE

46322 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46323 indicate the error.

## 46324 ERRORS

46325 The *tcgetattr()* function shall fail if:

46326 [EBADF] The *fildev* argument is not a valid file descriptor.

46327 [ENOTTY] The file associated with *fildev* is not a terminal.

## 46328 EXAMPLES

46329 None.

## 46330 APPLICATION USAGE

46331 None.

## 46332 RATIONALE

46333 Care must be taken when changing the terminal attributes. Applications should always do a  
46334 *tcgetattr()*, save the **termios** structure values returned, and then do a *tcsetattr()* changing only  
46335 the necessary fields. The application should use the values saved from the *tcgetattr()* to reset the  
46336 terminal state whenever it is done with the terminal. This is necessary because terminal  
46337 attributes apply to the underlying port and not to each individual open instance; that is, all  
46338 processes that have used the terminal see the latest attribute changes.

46339 A program that uses these functions should be written to catch all signals and take other  
46340 appropriate actions to ensure that when the program terminates, whether planned or not, the  
46341 terminal device's state is restored to its original state.

46342 Existing practice dealing with error returns when only part of a request can be honored is based  
46343 on calls to the *ioctl()* function. In historical BSD and System V implementations, the  
46344 corresponding *ioctl()* returns zero if the requested actions were semantically correct, even if  
46345 some of the requested changes could not be made. Many existing applications assume this  
46346 behavior and would no longer work correctly if the return value were changed from zero to -1  
46347 in this case.

46348 Note that either specification has a problem. When zero is returned, it implies everything  
46349 succeeded even if some of the changes were not made. When -1 is returned, it implies  
46350 everything failed even though some of the changes were made.

46351 Applications that need all of the requested changes made to work properly should follow  
46352 *tcsetattr()* with a call to *tcgetattr()* and compare the appropriate field values.

46353 **FUTURE DIRECTIONS**

46354 None.

46355 **SEE ALSO**

46356 *tcsetattr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**termios.h**>, the Base  
46357 Definitions volume of IEEE Std. 1003.1-200x, Chapter 11, General Terminal Interface

46358 **CHANGE HISTORY**

46359 First released in Issue 3.

46360 Entry included for alignment with the POSIX.1-1988 standard.

46361 **Issue 4**

46362 The FUTURE DIRECTIONS section is added to allow for alignment with the ISO POSIX-1  
46363 standard.

46364 **Issue 6**

46365 In the DESCRIPTION, the rate returned as the input baud rate shall be the output rate.  
46366 Previously, the number zero was also allowed but was obsolescent.

46367 **NAME**

46368 tcgetpgrp — get the foreground process group ID

46369 **SYNOPSIS**

46370 #include &lt;unistd.h&gt;

46371 pid\_t tcgetpgrp(int *fildev*);46372 **DESCRIPTION**46373 The *tcgetpgrp()* function shall return the value of the process group ID of the foreground process  
46374 group associated with the terminal.46375 If there is no foreground process group, *tcgetpgrp()* returns a value greater than 1 that does not  
46376 match the process group ID of any existing process group.46377 The *tcgetpgrp()* function is allowed from a process that is a member of a background process  
46378 group; however, the information may be subsequently changed by a process that is a member of  
46379 a foreground process group.46380 **RETURN VALUE**46381 Upon successful completion, *tcgetpgrp()* shall return the value of the process group ID of the  
46382 foreground process associated with the terminal. Otherwise, -1 shall be returned and *errno* set to  
46383 indicate the error.46384 **ERRORS**46385 The *tcgetpgrp()* function shall fail if:46386 [EBADF] The *fildev* argument is not a valid file descriptor. |46387 [ENOTTY] The calling process does not have a controlling terminal, or the file is not the |  
46388 controlling terminal.46389 **EXAMPLES**

46390 None.

46391 **APPLICATION USAGE**

46392 None.

46393 **RATIONALE**

46394 None.

46395 **FUTURE DIRECTIONS**

46396 None.

46397 **SEE ALSO**46398 *setsid()*, *setpgid()*, *tcsetpgrp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
46399 <sys/types.h>, <unistd.h>46400 **CHANGE HISTORY**

46401 First released in Issue 3.

46402 Entry included for alignment with the POSIX.1-1988 standard.

46403 **Issue 4**46404 The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
46405 XSI-conformant systems.

46406 The &lt;unistd.h&gt; header is added to the SYNOPSIS section.

46407 The following change is incorporated for alignment with the FIPS requirements:

- 46408           • The DESCRIPTION is clarified and the phrase “If \_POSIX\_JOB\_CONTROL is defined” is  
46409           removed because job control is now mandatory on all XSI-conformant systems.

46410 **Issue 6**

46411           In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

46412           The following new requirements on POSIX implementations derive from alignment with the  
46413           Single UNIX Specification:

- 46414           • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
46415           required for conforming implementations of previous POSIX specifications, it was not  
46416           required for UNIX applications.
- 46417           • In the DESCRIPTION, text previously conditional on support for \_POSIX\_JOB\_CONTROL is  
46418           now mandatory. This is a FIPS requirement.

46419 **NAME**

46420 tcgetsid — get process group ID for session leader for controlling terminal

46421 **SYNOPSIS**

46422 XSI #include &lt;termios.h&gt;

46423 pid\_t tcgetsid(int *fildes*);

46424

46425 **DESCRIPTION**46426 The *tcgetsid()* function shall obtain the process group ID of the session for which the terminal  
46427 specified by *fildes* is the controlling terminal.46428 **RETURN VALUE**46429 Upon successful completion, *tcgetsid()* shall return the process group ID associated with the  
46430 terminal. Otherwise, a value of (**pid\_t**)-1 shall be returned and *errno* set to indicate the error.46431 **ERRORS**46432 The *tcgetsid()* function shall fail if:46433 [EBADF] The *fildes* argument is not a valid file descriptor. |46434 [ENOTTY] The calling process does not have a controlling terminal, or the file is not the |  
46435 controlling terminal.46436 **EXAMPLES**

46437 None.

46438 **APPLICATION USAGE**

46439 None.

46440 **RATIONALE**

46441 None.

46442 **FUTURE DIRECTIONS**

46443 None.

46444 **SEE ALSO**

46445 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;termios.h&gt; |

46446 **CHANGE HISTORY**

46447 First released in Issue 4, Version 2.

46448 **Issue 5**

46449 Moved from X/OPEN UNIX extension to BASE.

46450 The [EACCES] error has been removed from the list of mandatory errors, and the description of  
46451 [ENOTTY] has been reworded.

46452 **NAME**

46453 tcsendbreak — send a “break” for a specific duration

46454 **SYNOPSIS**

46455 #include &lt;termios.h&gt;

46456 int tcsendbreak(int *fildev*, int *duration*);46457 **DESCRIPTION**46458 The *fildev* argument is an open file descriptor associated with a terminal.

46459 If the terminal is using asynchronous serial data transmission, *tcsendbreak()* shall cause  
46460 transmission of a continuous stream of zero-valued bits for a specific duration. If *duration* is 0, it  
46461 shall cause transmission of zero-valued bits for at least 0.25 seconds, and not more than 0.5  
46462 seconds. If *duration* is not 0, it shall send zero-valued bits for an implementation-defined period  
46463 of time.

46464 If the terminal is not using asynchronous serial data transmission, it is implementation-defined  
46465 whether *tcsendbreak()* sends data to generate a break condition or returns without taking any  
46466 action.

46467 Attempts to use *tcsendbreak()* from a process which is a member of a background process group  
46468 on a *fildev* associated with its controlling terminal, shall cause the process group to be sent a  
46469 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is  
46470 allowed to perform the operation, and no signal is sent.

46471 **RETURN VALUE**

46472 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46473 indicate the error.

46474 **ERRORS**46475 The *tcsendbreak()* function shall fail if:46476 [EBADF] The *fildev* argument is not a valid file descriptor.46477 [ENOTTY] The file associated with *fildev* is not a terminal.46478 The *tcsendbreak()* function may fail if:

46479 [EIO] The process group of the writing process is orphaned, and the writing process  
46480 is not ignoring or blocking SIGTTOU.

46481 **EXAMPLES**

46482 None.

46483 **APPLICATION USAGE**

46484 None.

46485 **RATIONALE**

46486 None.

46487 **FUTURE DIRECTIONS**

46488 None.

46489 **SEE ALSO**

46490 The Base Definitions volume of IEEE Std. 1003.1-200x, <termios.h>, <unistd.h>, the Base  
46491 Definitions volume of IEEE Std. 1003.1-200x, Chapter 11, General Terminal Interface

46492 **CHANGE HISTORY**

46493 First released in Issue 3.

46494 Entry included for alignment with the POSIX.1-1988 standard.

46495 **Issue 4**

46496 The [EIO] error is added to the ERRORS section.

46497 The following change is incorporated for alignment with the FIPS requirements:

- 46498
- In the DESCRIPTION the phrase “If \_POSIX\_JOB\_CONTROL is defined” is removed
- 46499 because job control is now mandatory on all XSI-conformant systems.

46500 **Issue 6**46501 The following new requirements on POSIX implementations derive from alignment with the  
46502 Single UNIX Specification:

- 46503
- In the DESCRIPTION, text previously conditional on \_POSIX\_JOB\_CONTROL is now
- 46504 mandated. This is a FIPS requirement.
- 
- 46505
- The [EIO] error is added.

## 46506 NAME

46507 tcsetattr — set the parameters associated with the terminal

46508 **Notes to Reviewers**46509 *This section with side shading will not appear in the final copy. - Ed.*

46510 See the FUTURE DIRECTIONS section.

46511 **SYNOPSIS**

46512 #include &lt;termios.h&gt;

46513 int tcsetattr(int *fildev*, int *optional\_actions*,46514 const struct termios \**termios\_p*);46515 **DESCRIPTION**46516 The *tcsetattr()* function shall set the parameters associated with the terminal referred to by the  
46517 open file descriptor *fildev* (an open file descriptor associated with a terminal) from the **termios**  
46518 structure referenced by *termios\_p* as follows:

- 46519 • If *optional\_actions* is TCSANOW, the change shall occur immediately.
- 46520 • If *optional\_actions* is TCSADRAIN, the change shall occur after all output written to *fildev* is  
46521 transmitted. This function should be used when changing parameters that affect output.
- 46522 • If *optional\_actions* is TCSAFLUSH, the change shall occur after all output written to *fildev* is  
46523 transmitted, and all input so far received but not read shall be discarded before the change is  
46524 made.

46525 If the output baud rate stored in the **termios** structure pointed to by *termios\_p* is the zero baud  
46526 rate, B0, the modem control lines shall no longer be asserted. Normally, this shall disconnect the  
46527 line.

46528 If the input baud rate stored in the **termios** structure pointed to by *termios\_p* is 0, the input baud  
46529 rate given to the hardware is the same as the output baud rate stored in the **termios** structure.

46530 The *tcsetattr()* function shall return successfully if it was able to perform any of the requested  
46531 actions, even if some of the requested actions could not be performed. It shall set all the  
46532 attributes that the implementation supports as requested and leaves all the attributes not  
46533 supported by the implementation unchanged. If no part of the request can be honored, it shall  
46534 return  $-1$  and set *errno* to [EINVAL]. If the input and output baud rates differ and are a  
46535 combination that is not supported, neither baud rate is changed. A subsequent call to *tcsetattr()*  
46536 shall return the actual state of the terminal device (reflecting both the changes made and not  
46537 made in the previous *tcsetattr()* call). The *tcsetattr()* function shall not change the values found  
46538 in the **termios** structure under any circumstances.

46539 The effect of *tcsetattr()* is undefined if the value of the **termios** structure pointed to by *termios\_p*  
46540 was not derived from the result of a call to *tcsetattr()* on *fildev*; an application should modify  
46541 only fields and flags defined by this volume of IEEE Std. 1003.1-200x between the call to  
46542 *tcsetattr()* and *tcsetattr()*, leaving all other fields and flags unmodified.

46543 No actions defined by this volume of IEEE Std. 1003.1-200x, other than a call to *tcsetattr()* or a  
46544 close of the last file descriptor in the system associated with this terminal device, shall cause any  
46545 of the terminal attributes defined by this volume of IEEE Std. 1003.1-200x to change.

46546 If *tcsetattr()* is called from a process which is a member of a background process group on a  
46547 *fildev* associated with its controlling terminal:

- 46548 • If the calling process is blocking or ignoring SIGTTOU signals, the operation completes  
46549 normally and no signal is sent.

46550           • Otherwise, a SIGTTOU signal shall be sent to the process group.

#### 46551 RETURN VALUE

46552           Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46553           indicate the error.

#### 46554 ERRORS

46555           The *tcsetattr()* function shall fail if:

46556           [EBADF]           The *fildev* argument is not a valid file descriptor.

46557           [EINTR]           A signal interrupted *tcsetattr()*.

46558           [EINVAL]          The *optional\_actions* argument is not a supported value, or an attempt was  
46559           made to change an attribute represented in the **termios** structure to an  
46560           unsupported value.

46561           [ENOTTY]          The file associated with *fildev* is not a terminal.

46562           The *tcsetattr()* function may fail if:

46563           [EIO]                The process group of the writing process is orphaned, and the writing process  
46564           is not ignoring or blocking SIGTTOU.

#### 46565 EXAMPLES

46566           None.

#### 46567 APPLICATION USAGE

46568           If trying to change baud rates, applications should call *tcsetattr()* then call *tcgetattr()* in order to  
46569           determine what baud rates were actually selected.

#### 46570 RATIONALE

46571           The *tcsetattr()* function can be interrupted in the following situations:

- 46572           • It is interrupted while waiting for output to drain.
- 46573           • It is called from a process in a background process group and SIGTTOU is caught.

46574           See also the RATIONALE section in *tcgetattr()*.

#### 46575 FUTURE DIRECTIONS

46576           Using an input baud rate of 0 to set the input rate equal to the output rate may not necessarily be  
46577           supported in a future version of this volume of IEEE Std. 1003.1-200x.

#### 46578 SEE ALSO

46579           *cfgetispeed()*, *tcgetattr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<termios.h>**,  
46580           **<unistd.h>**, the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 11, General Terminal  
46581           Interface

#### 46582 CHANGE HISTORY

46583           First released in Issue 3.

46584           Entry included for alignment with the POSIX.1-1988 standard.

#### 46585 Issue 4

46586           The words “and stores them in” are changed to “from” in the first paragraph of the  
46587           DESCRIPTION.

46588           The [EINTR] and [EIO] errors are added to the ERRORS section.

46589           The FUTURE DIRECTIONS section is added to allow for alignment with the ISO POSIX-1  
46590           standard.

- 46591 The following change is incorporated for alignment with the ISO POSIX-1 standard:
- 46592 • The argument *termios\_p* is changed from type **struct termios\*** to **const struct termios\***.
- 46593 The following change is incorporated for alignment with the FIPS requirements:
- 46594 • In the DESCRIPTION the phrase “If `_POSIX_JOB_CONTROL` is defined” is removed  
46595 because job control is now mandatory on all XSI-conformant systems.
- 46596 **Issue 6**
- 46597 The following new requirements on POSIX implementations derive from alignment with the  
46598 Single UNIX Specification:
- 46599 • In the DESCRIPTION, text previously conditional on `_POSIX_JOB_CONTROL` is now  
46600 mandated. This is a FIPS requirement.
- 46601 • The [EIO] error is added.
- 46602 In the DESCRIPTION, the text describing use of *tcsetattr()* from a process which is a member of  
46603 a background process group is clarified.

46604 **NAME**

46605 tcsetpgrp — set the foreground process group ID

46606 **SYNOPSIS**

46607 #include &lt;unistd.h&gt;

46608 int tcsetpgrp(int *fildev*, pid\_t *pgid\_id*);46609 **DESCRIPTION**

46610 If the process has a controlling terminal, *tcsetpgrp()* shall set the foreground process group ID  
46611 associated with the terminal to *pgid\_id*. The application shall ensure that the file associated with  
46612 *fildev* is the controlling terminal of the calling process and the controlling terminal is currently  
46613 associated with the session of the calling process. The application shall ensure that the value of  
46614 *pgid\_id* matches a process group ID of a process in the same session as the calling process.

46615 Attempts to use *tcsetpgrp()* from a process which is a member of a background process group on  
46616 a *fildev* associated with its controlling terminal will cause the process group to be sent a  
46617 SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is  
46618 allowed to perform the operation, and no signal is sent.

46619 **RETURN VALUE**

46620 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to  
46621 indicate the error.

46622 **ERRORS**46623 The *tcsetpgrp()* function shall fail if:

46624 [EBADF] The *fildev* argument is not a valid file descriptor.

46625 [EINVAL] This implementation does not support the value in the *pgid\_id* argument.

46626 [ENOTTY] The calling process does not have a controlling terminal, or the file is not the  
46627 controlling terminal, or the controlling terminal is no longer associated with  
46628 the session of the calling process.

46629 [EPERM] The value of *pgid\_id* is a value supported by the implementation, but does not  
46630 match the process group ID of a process in the same session as the calling  
46631 process.

46632 **EXAMPLES**

46633 None.

46634 **APPLICATION USAGE**

46635 None.

46636 **RATIONALE**

46637 None.

46638 **FUTURE DIRECTIONS**

46639 None.

46640 **SEE ALSO**46641 *tcgetpgrp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <unistd.h>46642 **CHANGE HISTORY**

46643 First released in Issue 3.

46644 Entry included for alignment with the POSIX.1-1988 standard.

46645 **Issue 4**

46646 The `<sys/types.h>` header is now marked as optional (OH); this header need not be included on  
46647 XSI-conformant systems.

46648 The `<unistd.h>` header is added to the SYNOPSIS section.

46649 The [ENOSYS] error is removed from the ERRORS section.

46650 The following change is incorporated for alignment with the FIPS requirements:

- 46651 • In the DESCRIPTION the phrase “If \_POSIX\_JOB\_CONTROL is defined” is removed  
46652 because job control is now mandatory on all XSI-conformant systems.

46653 **Issue 6**

46654 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

46655 The following new requirements on POSIX implementations derive from alignment with the  
46656 Single UNIX Specification:

- 46657 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
46658 required for conforming implementations of previous POSIX specifications, it was not  
46659 required for UNIX applications.
- 46660 • In the DESCRIPTION and ERRORS sections, text previously conditional on  
46661 `_POSIX_JOB_CONTROL` is now mandated. This is a FIPS requirement.

46662 The DESCRIPTION is updated to avoid use of the term “must” for application requirements. |

46663 The Open Group corrigenda item U047/4 has been applied. |

## 46664 NAME

46665 tdelete, tfind, tsearch, twalk — manage a binary search tree

## 46666 SYNOPSIS

```
46667 xsi #include <search.h>
46668 void *tdelete(const void *restrict key, void **restrict rootp,
46669 int(*compar)(const void *, const void *));
46670 void *tfind(const void *key, void *const *rootp,
46671 int(*compar)(const void *, const void *));
46672 void *tsearch(const void *key, void **rootp,
46673 int (*compar)(const void *, const void *));
46674 void twalk(const void *root,
46675 void (*action)(const void *, VISIT, int));
46676
```

## 46677 DESCRIPTION

46678 The *tdelete()*, *tfind()*, *tsearch()*, and *twalk()* functions manipulate binary search trees.  
 46679 Comparisons are made with a user-supplied routine, the address of which is passed as the  
 46680 *compar* argument. This routine is called with two arguments, the pointers to the elements being  
 46681 compared. The application shall ensure that the user-supplied routine returns an integer less  
 46682 than, equal to, or greater than 0, according to whether the first argument is to be considered less  
 46683 than, equal to, or greater than the second argument. The comparison function need not compare  
 46684 every byte, so arbitrary data may be contained in the elements in addition to the values being  
 46685 compared.

46686 The *tsearch()* function is used to build and access the tree. The *key* argument is a pointer to an  
 46687 element to be accessed or stored. If there is a node in the tree whose element is equal to the value  
 46688 pointed to by *key*, a pointer to this found node is returned. Otherwise, the value pointed to by  
 46689 *key* is inserted (that is, a new node is created and the value of *key* is copied to this node), and a  
 46690 pointer to this node returned. Only pointers are copied, so the application shall ensure that the  
 46691 calling routine stores the data. The *rootp* argument points to a variable that points to the root  
 46692 node of the tree. A null pointer value for the variable pointed to by *rootp* denotes an empty tree;  
 46693 in this case, the variable shall be set to point to the node which shall be at the root of the new  
 46694 tree.

46695 Like *tsearch()*, *tfind()* shall search for a node in the tree, returning a pointer to it if found.  
 46696 However, if it is not found, *tfind()* shall return a null pointer. The arguments for *tfind()* are the  
 46697 same as for *tsearch()*.

46698 The *tdelete()* function deletes a node from a binary search tree. The arguments are the same as  
 46699 for *tsearch()*. The variable pointed to by *rootp* shall be changed if the deleted node was the root  
 46700 of the tree. The *tdelete()* function returns a pointer to the parent of the deleted node, or a null  
 46701 pointer if the node is not found.

46702 The *twalk()* function traverses a binary search tree. The *root* argument is a pointer to the root  
 46703 node of the tree to be traversed. (Any node in a tree may be used as the root for a walk below  
 46704 that node.) The argument *action* is the name of a routine to be invoked at each node. This routine  
 46705 is, in turn, called with three arguments. The first argument is the address of the node being  
 46706 visited. The structure pointed to by this argument is unspecified and shall not be modified by  
 46707 the application, but it is guaranteed that a pointer-to-node can be converted to pointer-to-  
 46708 pointer-to-element to access the element stored in the node. The second argument is a value  
 46709 from an enumeration data type:

```
46710 typedef enum { preorder, postorder, endorder, leaf } VISIT;
```

46711 (defined in `<search.h>`), depending on whether this is the first, second, or third time that the  
 46712 node is visited (during a depth-first, left-to-right traversal of the tree), or whether the node is a  
 46713 leaf. The third argument is the level of the node in the tree, with the root being level 0.

46714 If the calling function alters the pointer to the root, the result is undefined.

#### 46715 RETURN VALUE

46716 If the node is found, both `tsearch()` and `tfind()` shall return a pointer to it. If not, `tfind()` shall  
 46717 return a null pointer, and `tsearch()` shall return a pointer to the inserted item.

46718 A null pointer shall be returned by `tsearch()` if there is not enough space available to create a new  
 46719 node.

46720 A null pointer shall be returned by `tdelete()`, `tfind()`, and `tsearch()` if `rootp` is a null pointer on  
 46721 entry.

46722 The `tdelete()` function shall return a pointer to the parent of the deleted node, or a null pointer if  
 46723 the node is not found.

46724 The `twalk()` function shall return no value.

#### 46725 ERRORS

46726 No errors are defined.

#### 46727 EXAMPLES

46728 The following code reads in strings and stores structures containing a pointer to each string and  
 46729 a count of its length. It then walks the tree, printing out the stored strings and their lengths in  
 46730 alphabetical order.

```

46731 #include <search.h>
46732 #include <string.h>
46733 #include <stdio.h>

46734 #define STRSZ 10000
46735 #define NODSZ 500

46736 struct node { /* Pointers to these are stored in the tree. */
46737 char *string;
46738 int length;
46739 };

46740 char string_space[STRSZ]; /* Space to store strings. */
46741 struct node nodes[NODSZ]; /* Nodes to store. */
46742 void *root = NULL; /* This points to the root. */

46743 int main(int argc, char *argv[])
46744 {
46745 char *strptr = string_space;
46746 struct node *nodeptr = nodes;
46747 void print_node(const void *, VISIT, int);
46748 int i = 0, node_compare(const void *, const void *);

46749 while (gets(strptr) != NULL && i++ < NODSZ) {
46750 /* Set node. */
46751 nodeptr->string = strptr;
46752 nodeptr->length = strlen(strptr);
46753 /* Put node into the tree. */
46754 (void) tsearch((void *)nodeptr, (void **)&root,
46755 node_compare);

```

```

46756 /* Adjust pointers, so we do not overwrite tree. */
46757 strptr += nodeptr->length + 1;
46758 nodeptr++;
46759 }
46760 twalk(root, print_node);
46761 return 0;
46762 }
46763 /*
46764 * This routine compares two nodes, based on an
46765 * alphabetical ordering of the string field.
46766 */
46767 int
46768 node_compare(const void *node1, const void *node2)
46769 {
46770 return strcmp(((const struct node *) node1)->string,
46771 ((const struct node *) node2)->string);
46772 }
46773 /*
46774 * This routine prints out a node, the second time
46775 * twalk encounters it or if it is a leaf.
46776 */
46777 void
46778 print_node(const void *ptr, VISIT order, int level)
46779 {
46780 const struct node *p = *(const struct node **) ptr;
46781 if (order == postorder || order == leaf) {
46782 (void) printf("string = %s, length = %d\n",
46783 p->string, p->length);
46784 }
46785 }

```

#### 46786 APPLICATION USAGE

46787 The *root* argument to *twalk()* is one level of indirection less than the *rootp* arguments to *tdelete()* and *tsearch()*.

46789 There are two nomenclatures used to refer to the order in which tree nodes are visited. The *tsearch()* function uses **preorder**, **postorder**, and **endorder** to refer respectively to visiting a node before any of its children, after its left child and before its right, and after both its children. The alternative nomenclature uses **preorder**, **inorder**, and **postorder** to refer to the same visits, which could result in some confusion over the meaning of **postorder**.

#### 46794 RATIONALE

46795 None.

#### 46796 FUTURE DIRECTIONS

46797 None.

#### 46798 SEE ALSO

46799 *hcreate()*, *tsearch()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <[search.h](#)>

46800 **CHANGE HISTORY**

- 46801 First released in Issue 1. Derived from Issue 1 of the SVID.
- 46802 **Issue 4**
- 46803 The type of argument *key* is changed from **char\*** to **const void\***.
- 46804 The function return value is changed from **char\*** to **void\***.
- 46805 Arguments to *compar* are formally defined.
- 46806 The type of argument *rootp* is changed from **char\*\*** to **void\*\*** for the *tsearch()* function.
- 46807 The type of argument *rootp* is changed from **char\*\*** to **void\*const\*** for the *tfind()* function.
- 46808 The type of argument *root* is changed from **char\*** to **const void\***, and the argument list to *action* is formally defined for the *twalk()* function.
- 46809
- 46810 Various minor wording changes are made in the DESCRIPTION to improve clarity and accuracy. In particular, additional notes are added about constraints on the first argument to *twalk()*.
- 46811
- 46812
- 46813 The sample code in the EXAMPLES section is updated to use ISO C standard syntax. Also the definition of the *root* and *argv* items is changed.
- 46814
- 46815 The paragraph in the APPLICATION USAGE section about casts is removed.
- 46816 **Issue 5**
- 46817 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in previous issues.
- 46818
- 46819 **Issue 6**
- 46820 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.
- 46821 The **restrict** keyword is added to the *tdelete()* prototype for alignment with the ISO/IEC 9899:1999 standard.
- 46822

46823 **NAME**

46824 telfdir — current location of a named directory stream

46825 **SYNOPSIS**

46826 XSI #include &lt;dirent.h&gt;

46827 long telfdir(DIR \*dirp);

46828

46829 **DESCRIPTION**46830 The *telfdir()* function obtains the current location associated with the directory stream specified  
46831 by *dirp*.46832 If the most recent operation on the directory stream was a *seekdir()*, the directory position  
46833 returned from the *telfdir()* shall be the same as that supplied as a *loc* argument for *seekdir()*.46834 **RETURN VALUE**46835 Upon successful completion, *telfdir()* shall return the current location of the specified directory  
46836 stream.46837 **ERRORS**

46838 No errors are defined.

46839 **EXAMPLES**

46840 None.

46841 **APPLICATION USAGE**

46842 None.

46843 **RATIONALE**

46844 None.

46845 **FUTURE DIRECTIONS**

46846 None.

46847 **SEE ALSO**46848 *opendir()*, *readdir()*, *seekdir()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <dirent.h>46849 **CHANGE HISTORY**

46850 First released in Issue 2.

46851 **Issue 4**

46852 The &lt;sys/types.h&gt; header is removed from the SYNOPSIS section.

46853 The function return value is expanded to **long**.46854 **Issue 4, Version 2**46855 The DESCRIPTION is updated for X/OPEN UNIX conformance to indicate that a call to *telfdir()*  
46856 immediately following a call to *seekdir()*, returns the *loc* value passed to the *seekdir()* call.

46857 **NAME**

46858 tempnam — create a name for a temporary file

46859 **SYNOPSIS**46860 XSI 

```
#include <stdio.h>
```

46861 

```
char *tempnam(const char *dir, const char *pfx);
```

46862

46863 **DESCRIPTION**46864 The *tempnam()* function generates a path name that may be used for a temporary file.

46865 The *tempnam()* function allows the user to control the choice of a directory. The *dir* argument  
46866 points to the name of the directory in which the file is to be created. If *dir* is a null pointer or  
46867 points to a string which is not a name for an appropriate directory, the path prefix defined as  
46868 P\_tmpdir in the <stdio.h> header is used. If that directory is not accessible, an implementation-  
46869 defined directory may be used.

46870 Many applications prefer their temporary files to have certain initial letter sequences in their  
46871 names. The *pfx* argument should be used for this. This argument may be a null pointer or point  
46872 to a string of up to five bytes to be used as the beginning of the file name.

46873 Some implementations of *tempnam()* may use *tmpnam()* internally. On such implementations, if  
46874 called more than {TMP\_MAX} times in a single process, the behavior is implementation-defined.

46875 **RETURN VALUE**

46876 Upon successful completion, *tempnam()* shall allocate space for a string, put the generated path  
46877 name in that space, and return a pointer to it. The pointer shall be suitable for use in a  
46878 subsequent call to *free()*. Otherwise, it shall return a null pointer and set *errno* to indicate the  
46879 error.

46880 **ERRORS**46881 The *tempnam()* function shall fail if:

46882 [ENOMEM] Insufficient storage space is available.

46883 **EXAMPLES**46884 **Generating a Path Name**

46885 The following example generates a path name for a temporary file in directory */tmp*, with the  
46886 prefix *file*. After the file name has been created, the call to *free()* deallocates the space used to  
46887 store the file name.

```
46888 #include <stdio.h>
46889 #include <stdlib.h>
46890 ...
46891 char *directory = "/tmp";
46892 char *fileprefix = "file";
46893 char *file;

46894 file = tempnam(directory, fileprefix);
46895 free(file);
```

46896 **APPLICATION USAGE**

46897 This function only creates path names. It is the application's responsibility to create and remove  
46898 the files. Between the time a path name is created and the file is opened, it is possible for some  
46899 other process to create a file with the same name. Applications may find *tmpfile()* more useful.

46900 **RATIONALE**

46901 None.

46902 **FUTURE DIRECTIONS**

46903 None.

46904 **SEE ALSO**

46905 *fopen()*, *free()*, *open()*, *tmpfile()*, *tmpnam()*, *unlink()*, the Base Definitions volume of  
46906 IEEE Std. 1003.1-200x, <stdio.h>

46907 **CHANGE HISTORY**

46908 First released in Issue 1. Derived from Issue 1 of the SVID.

46909 **Issue 4**

46910 The type of arguments *dir* and *px* is changed from **char\*** to **const char\***.

46911 The DESCRIPTION is changed to indicate that *px* is treated as a string of bytes and not as a  
46912 string of (possibly multi-byte) characters.

46913 The second paragraph of the APPLICATION USAGE section is expanded.

46914 **Issue 5**

46915 The last paragraph of the DESCRIPTION was included as an APPLICATION USAGE note in  
46916 previous issues.

46917 **NAME**

46918 tfind — search binary search tree

46919 **SYNOPSIS**

46920 XSI #include &lt;search.h&gt;

46921 void \*tfind(const void \*key, void \*const \*rootp,  
46922 int (\*compar)(const void \*, const void \*));

46923

46924 **DESCRIPTION**46925 Refer to *tdelete()*.

46926 **NAME**

46927 tgamma, tgammaf, tgammal — compute gamma() function

46928 **SYNOPSIS**

46929 #include <math.h>

46930 double tgamma(double x);

46931 float tgammaf(float x);

46932 long double tgammal(long double x);

46933 **DESCRIPTION**

46934 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
46935 conflict between the requirements described here and the ISO C standard is unintentional. This  
46936 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

46937 These functions shall compute the *gamma()* function of *x*.

46938 An application wishing to check for error situations should set *errno* to 0 before calling these  
46939 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

46940 **RETURN VALUE**

46941 Upon successful completion, these functions shall return *Gamma(x)*.

46942 If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

46943 If *x* is a negative integer, or if the result cannot be represented when *x* is 0, either HUGE\_VAL or  
46944 NaN shall be returned and *errno* shall be set to [EDOM].

46945 If the magnitude of *x* is too large or too small, ±HUGE\_VAL shall be returned and *errno* shall be  
46946 set to [ERANGE].

46947 **ERRORS**

46948 These functions shall fail if:

46949 [EDOM] The value of *x* is negative or the result cannot be represented when *x* is zero.

46950 [ERANGE] The magnitude of *x* is too large or too small.

46951 These functions may fail if:

46952 [EDOM] The value of *x* is NaN.

46953 **EXAMPLES**

46954 None.

46955 **APPLICATION USAGE**

46956 None.

46957 **RATIONALE**

46958 This function is named *tgamma()* in order to avoid conflicts with the historical *gamma()* and  
46959 *lgamma()* functions.

46960 **FUTURE DIRECTIONS**

46961 None.

46962 **SEE ALSO**

46963 *lgamma()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

46964 **CHANGE HISTORY**

46965 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

46966 **NAME**

46967           time — get time

46968 **SYNOPSIS**

46969           #include &lt;time.h&gt;

46970           time\_t time(time\_t \*tloc);

46971 **DESCRIPTION**

46972 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
 46973 conflict between the requirements described here and the ISO C standard is unintentional. This  
 46974 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

46975 **CX**       The *time()* function returns the value of time in seconds since the Epoch.

46976           The *tloc* argument points to an area where the return value is also stored. If *tloc* is a null pointer,  
 46977 no value is stored.

46978 **RETURN VALUE**

46979           Upon successful completion, *time()* shall return the value of time. Otherwise, **(time\_t)−1** shall be  
 46980 returned.

46981 **ERRORS**

46982           No errors are defined.

46983 **EXAMPLES**46984           **Getting the Current Time**

46985           The following example uses the *time()* function to calculate the time elapsed, in seconds, since  
 46986 January 1, 1970 0:00 UTC, *localtime()* to convert that value to a broken-down time, and *asctime()*  
 46987 to convert the broken-down time values into a printable string.

46988           #include &lt;stdio.h&gt;

46989           #include &lt;time.h&gt;

46990           main()

46991           {

46992           time\_t result;

46993                 result = time(NULL);

46994                 printf("%s%ld secs since the Epoch\n",

46995                         asctime(localtime(&amp;result)),

46996                         (long)result);

46997                 return(0);

46998           }

46999           This example writes the current time to *stdout* in a form like this:

47000           Wed Jun 26 10:32:15 1996

47001           835810335 secs since the Epoch

47002 **Timing an Event**

47003 The following example gets the current time, prints it out in the user's format, and prints the  
47004 number of minutes to an event being timed.

```
47005 #include <time.h>
47006 #include <stdio.h>
47007 ...
47008 time_t now;
47009 int minutes_to_event;
47010 ...
47011 time(&now);
47012 printf("The time is ");
47013 puts(asctime(localtime(&now)));
47014 printf("There are %d minutes to the event.\n",
47015 minutes_to_event);
47016 ...
```

47017 **APPLICATION USAGE**

47018 None.

47019 **RATIONALE**

47020 The *time()* function returns a value in seconds (type **time\_t**) while *times()* returns a set of values  
47021 in clock ticks (type **clock\_t**). Some historical implementations, such as 4.3 BSD, have  
47022 mechanisms capable of returning more precise times (see below). A generalized timing scheme  
47023 to unify these various timing mechanisms has been proposed but not adopted.

47024 Implementations in which **time\_t** is a 32-bit signed integer (many historical implementations)  
47025 fail in the year 2038. This issue of this volume of IEEE Std. 1003.1-200x does not address this  
47026 problem. However, the use of the **time\_t** type is mandated in order to ease the eventual fix.

47027 The use of the **<time.h>**, header instead of **<sys/types.h>**, allows compatibility with the ISO C  
47028 standard.

47029 Many historical implementations (including Version 7) and the 1984 /usr/group standard use  
47030 **long** instead of **time\_t**. This volume of IEEE Std. 1003.1-200x uses the latter type in order to  
47031 agree with the ISO C standard.

47032 4.3 BSD includes *time()* only as an alternate function to the more flexible *gettimeofday()* function.

47033 **FUTURE DIRECTIONS**

47034 In a future version of this volume of IEEE Std. 1003.1-200x, **time\_t** is likely to be required to be  
47035 capable of representing times far in the future. Whether this will be mandated as a 64-bit type or  
47036 a requirement that a specific date in the future be representable (for example, 10000 AD) is not  
47037 yet determined. Systems purchased after the approval of this volume of IEEE Std. 1003.1-200x  
47038 should be evaluated to determine whether their lifetime will extend past 2038.

47039 **SEE ALSO**

47040 *asctime()*, *clock()*, *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *mktime()*, *strftime()*, *strptime()*, *utime()*,  
47041 the Base Definitions volume of IEEE Std. 1003.1-200x, **<time.h>**

47042 **CHANGE HISTORY**

47043 First released in Issue 1. Derived from Issue 1 of the SVID.

47044 **Issue 4**

47045 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 47046 • The RETURN VALUE section is updated to indicate that (**time\_t**)−1 is returned on error.

47047 **Issue 6**

47048 Extensions beyond the ISO C standard are now marked.

## 47049 NAME

47050 timer\_create — create a per-process timer (**REALTIME**)

## 47051 SYNOPSIS

47052 TMR #include &lt;signal.h&gt;

47053 #include &lt;time.h&gt;

47054 int timer\_create(clockid\_t clockid, struct sigevent \*restrict evp,

47055 timer\_t \*restrict timerid);

47056

## 47057 DESCRIPTION

47058 The *timer\_create()* function shall create a per-process timer using the specified clock, *clock\_id*, as  
 47059 the timing base. The *timer\_create()* function returns, in the location referenced by *timerid*, a timer  
 47060 ID of type **timer\_t** used to identify the timer in timer requests. This timer ID shall be unique  
 47061 within the calling process until the timer is deleted. The particular clock, *clock\_id*, is defined in  
 47062 <**time.h**>. The timer whose ID is returned shall be in a disarmed state upon return from  
 47063 *timer\_create()*.

47064 The *evp* argument, if non-NULL, points to a **sigevent** structure. This structure, allocated by the  
 47065 application, defines the asynchronous notification to occur as specified in Section 2.4.1 (on page  
 47066 528) when the timer expires. If the *evp* argument is NULL, the effect is as if the *evp* argument  
 47067 pointed to a **sigevent** structure with the *sigev\_notify* member having the value SIGEV\_SIGNAL,  
 47068 the *sigev\_signo* having a default signal number, and the *sigev\_value* member having the value of  
 47069 the timer ID.

47070 Each implementation shall define a set of clocks that can be used as timing bases for per-process  
 47071 MON timers. All implementations shall support a *clock\_id* of CLOCK\_REALTIME. If the Monotonic  
 47072 Clock option is supported, implementations shall support a *clock\_id* of CLOCK\_MONOTONIC.

47073 Per-process timers shall not be inherited by a child process across a *fork()* and shall be disarmed  
 47074 and deleted by an *exec*.

47075 CPT If \_POSIX\_CPUTIME is defined, implementations shall support *clock\_id* values representing the  
 47076 CPU-time clock of the calling process.

47077 TCT If \_POSIX\_THREAD\_CPUTIME is defined, implementations shall support *clock\_id* values  
 47078 representing the CPU-time clock of the calling thread.

47079 CPT|TCT It is implementation-defined whether a *timer\_create()* function will succeed if the value defined  
 47080 by *clock\_id* corresponds to the CPU-time clock of a process or thread different from the process  
 47081 or thread invoking the function.

## 47082 RETURN VALUE

47083 If the call succeeds, *timer\_create()* shall return zero and update the location referenced by *timerid*  
 47084 to a **timer\_t**, which can be passed to the per-process timer calls. If an error occurs, the function  
 47085 shall return a value of -1 and set *errno* to indicate the error. The value of *timerid* is undefined if  
 47086 an error occurs.

## 47087 ERRORS

47088 The *timer\_create()* function shall fail if:

47089 [EAGAIN] The system lacks sufficient signal queuing resources to honor the request.

47090 [EAGAIN] The calling process has already created all of the timers it is allowed by this  
47091 implementation.

47092 [EINVAL] The specified clock ID is not defined.

47093 CPT|TCT [ENOTSUP] The implementation does not support the creation of a timer attached to the  
 47094 CPU-time clock that is specified by *clock\_id* and associated with a process or  
 47095 thread different from the process or thread invoking *timer\_create()*.

#### 47096 EXAMPLES

47097 None.

#### 47098 APPLICATION USAGE

47099 None.

#### 47100 RATIONALE

##### 47101 **Periodic Timer Overrun and Resource Allocation**

47102 The specified timer facilities may deliver realtime signals (that is, queued signals) on  
 47103 implementations that support this option. Because realtime applications cannot afford to lose  
 47104 notifications of asynchronous events, like timer expirations or asynchronous I/O completions, it  
 47105 must be possible to ensure that sufficient resources exist to deliver the signal when the event  
 47106 occurs. In general, this is not a difficulty because there is a one-to-one correspondence between a  
 47107 request and a subsequent signal generation. If the request cannot allocate the signal delivery  
 47108 resources, it can fail the call with an [EAGAIN] error.

47109 Periodic timers are a special case. A single request can generate an indeterminate number of  
 47110 signals. This is not a problem if the requesting process can service the signals as fast as they are  
 47111 generated, thus making the signal delivery resources available for delivery of subsequent  
 47112 periodic timer expiration signals. But, in general, this cannot be assured—processing of periodic  
 47113 timer signals may “overrun”; that is, subsequent periodic timer expirations may occur before the  
 47114 currently pending signal has been delivered.

47115 Also, for signals, according to the POSIX.1-1990 standard, if subsequent occurrences of a  
 47116 pending signal are generated, it is implementation-defined whether a signal is delivered for each  
 47117 occurrence. This is not adequate for some realtime applications. So a mechanism is required to  
 47118 allow applications to detect how many timer expirations were delayed without requiring an  
 47119 indefinite amount of system resources to store the delayed expirations.

47120 The specified facilities provide for an overrun count. The overrun count is defined as the number  
 47121 of extra timer expirations that occurred between the time a timer expiration signal is generated  
 47122 and the time the signal is delivered. The signal-catching function, if it is concerned with  
 47123 overruns, can retrieve this count on entry. With this method, a periodic timer only needs one  
 47124 “signal queuing resource” that can be allocated at the time of the *timer\_create()* function call.

47125 A function is defined to retrieve the overrun count so that an application need not allocate static  
 47126 storage to contain the count, and an implementation need not update this storage  
 47127 asynchronously on timer expirations. But, for some high-frequency periodic applications, the  
 47128 overhead of an additional system call on each timer expiration may be prohibitive. The  
 47129 functions, as defined, permit an implementation to maintain the overrun count in user space,  
 47130 associated with the *timerid*. The *timer\_getoverrun()* function can then be implemented as a macro  
 47131 that uses the *timerid* argument (which may just be a pointer to a user space structure containing  
 47132 the counter) to locate the overrun count with no system call overhead. Other implementations,  
 47133 less concerned with this class of applications, can avoid the asynchronous update of user space  
 47134 by maintaining the count in a system structure at the cost of the extra system call to obtain it.

47135 **Timer Expiration Signal Parameters**

47136 The Realtime Signals Extension option supports an application-specific datum that is delivered  
47137 to the extended signal handler. This value is explicitly specified by the application, along with  
47138 the signal number to be delivered, in a **sigevent** structure. The type of the application-defined  
47139 value can be either an integer constant or a pointer. This explicit specification of the value, as  
47140 opposed to always sending the timer ID, was selected based on existing practice.

47141 It is common practice for realtime applications (on non-POSIX systems or realtime extended  
47142 POSIX systems) to use the parameters of event handlers as the case label of a switch statement  
47143 or as a pointer to an application-defined data structure. Because *timer\_ids* are dynamically  
47144 allocated by the *timer\_create()* function, they can be used for neither of these functions without  
47145 additional application overhead in the signal handler; for example, to search an array of saved  
47146 timer IDs to associate the ID with a constant or application data structure.

47147 **FUTURE DIRECTIONS**

47148 None.

47149 **SEE ALSO**

47150 *clock\_getres()*, *timer\_delete()*, *timer\_getoverrun()*, the Base Definitions volume of  
47151 IEEE Std. 1003.1-200x, <**time.h**>

47152 **CHANGE HISTORY**

47153 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

47154 **Issue 6**

47155 The *timer\_create()* function is marked as part of the Timers option.

47156 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
47157 implementation does not support the Timers option.

47158 CPU-time clocks are added for alignment with IEEE Std. 1003.1d-1999.

47159 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by adding the  
47160 requirement for the CLOCK\_MONOTONIC clock under the Monotonic Clock option.

47161 The **restrict** keyword is added to the *timer\_create()* prototype for alignment with the  
47162 ISO/IEC 9899:1999 standard.

47163 **NAME**

47164 timer\_delete — delete a per-process timer (**REALTIME**)

47165 **SYNOPSIS**

```
47166 TMR #include <time.h>
```

```
47167 int timer_delete(timer_t timerid);
```

47168

47169 **DESCRIPTION**

47170 The *timer\_delete()* function deletes the specified timer, *timerid*, previously created by the  
47171 *timer\_create()* function. If the timer is armed when *timer\_delete()* is called, the behavior shall be  
47172 as if the timer is automatically disarmed before removal. The disposition of pending signals for  
47173 the deleted timer is unspecified.

47174 **RETURN VALUE**

47175 If successful, the *timer\_delete()* function shall return a value of zero. Otherwise, the function shall  
47176 return a value of  $-1$  and set *errno* to indicate the error.

47177 **ERRORS**

47178 The *timer\_delete()* function shall fail if:

47179 [EINVAL] The timer ID specified by *timerid* is not a valid timer ID.

47180 **EXAMPLES**

47181 None.

47182 **APPLICATION USAGE**

47183 None.

47184 **RATIONALE**

47185 None.

47186 **FUTURE DIRECTIONS**

47187 None.

47188 **SEE ALSO**

47189 *timer\_create()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>

47190 **CHANGE HISTORY**

47191 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

47192 **Issue 6**

47193 The *timer\_delete()* function is marked as part of the Timers option.

47194 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
47195 implementation does not support the Timers option.

## 47196 NAME

47197 timer\_getoverrun, timer\_gettime, timer\_settime — per-process timers (**REALTIME**)

## 47198 SYNOPSIS

47199 TMR #include &lt;time.h&gt;

```

47200 int timer_getoverrun(timer_t timerid);
47201 int timer_gettime(timer_t timerid, struct itimerspec *value);
47202 int timer_settime(timer_t timerid, int flags,
47203 const struct itimerspec *restrict value,
47204 struct itimerspec *restrict ovalue);
47205

```

## 47206 DESCRIPTION

47207 Only a single signal shall be queued to the process for a given timer at any point in time. When a  
 47208 timer for which a signal is still pending expires, no signal shall be queued, and a timer overrun  
 47209 RTS shall occur. When a timer expiration signal is delivered to or accepted by a process, if the  
 47210 implementation supports the Realtime Signals Extension, the *timer\_getoverrun()* function returns  
 47211 the timer expiration overrun count for the specified timer. The overrun count returned contains  
 47212 the number of extra timer expirations that occurred between the time the signal was generated  
 47213 (queued) and when it was delivered or accepted, up to but not including an implementation-  
 47214 defined maximum of {DELAYTIMER\_MAX}. If the number of such extra expirations is greater  
 47215 than or equal to {DELAYTIMER\_MAX}, then the overrun count is set to {DELAYTIMER\_MAX}.  
 47216 The value returned by *timer\_getoverrun()* applies to the most recent expiration signal delivery or  
 47217 acceptance for the timer. If no expiration signal has been delivered for the timer, or if the  
 47218 Realtime Signals Extension is not supported, the return value of *timer\_getoverrun()* is  
 47219 unspecified.

47220 The *timer\_gettime()* function shall store the amount of time until the specified timer, *timerid*,  
 47221 expires and the reload value of the timer into the space pointed to by the *value* argument. The  
 47222 *it\_value* member of this structure shall contain the amount of time before the timer expires, or  
 47223 zero if the timer is disarmed. This value is returned as the interval until timer expiration, even if  
 47224 the timer was armed with absolute time. The *it\_interval* member of *value* shall contain the reload  
 47225 value last set by *timer\_settime()*.

47226 The *timer\_settime()* function shall set the time until the next expiration of the timer specified by  
 47227 *timerid* from the *it\_value* member of the *value* argument and arms the timer if the *it\_value*  
 47228 member of *value* is non-zero. If the specified timer was already armed when *timer\_settime()* is  
 47229 called, this call shall reset the time until next expiration to the *value* specified. If the *it\_value*  
 47230 member of *value* is zero, the timer shall be disarmed. The effect of disarming or resetting a timer  
 47231 on pending expiration notifications is unspecified.

47232 If the flag **TIMER\_ABSTIME** is not set in the argument *flags*, *timer\_settime()* behaves as if the  
 47233 time until next expiration is set to be equal to the interval specified by the *it\_value* member of  
 47234 *value*. That is, the timer shall expire in *it\_value* nanoseconds from when the call is made. If the  
 47235 flag **TIMER\_ABSTIME** is set in the argument *flags*, *timer\_settime()* behaves as if the time until  
 47236 next expiration is set to be equal to the difference between the absolute time specified by the  
 47237 *it\_value* member of *value* and the current value of the clock associated with *timerid*. That is, the  
 47238 timer shall expire when the clock reaches the value specified by the *it\_value* member of *value*. If  
 47239 the specified time has already passed, the function shall succeed and the expiration notification  
 47240 shall be made.

47241 The reload value of the timer is set to the value specified by the *it\_interval* member of *value*.  
 47242 When a timer is armed with a non-zero *it\_interval*, a periodic (or repetitive) timer is specified.

47243 Time values that are between two consecutive non-negative integer multiples of the resolution  
 47244 of the specified timer shall be rounded up to the larger multiple of the resolution. Quantization  
 47245 error shall not cause the timer to expire earlier than the rounded time value.

47246 If the argument *ovalue* is not NULL, the function *timer\_settime()* shall store, in the location  
 47247 referenced by *ovalue*, a value representing the previous amount of time before the timer would  
 47248 have expired, or zero if the timer was disarmed, together with the previous timer reload value.

#### 47249 **Notes to Reviewers**

47250 *This section with side shading will not appear in the final copy. - Ed.*

47251 D1, XSH, ERN 388 suggests rewording “are subject to the resolution of the timer” to “may have  
 47252 been rounded to the resolution of the timer”.

47253 The members of *ovalue* are subject to the resolution of the timer, and they are the same values  
 47254 that would be returned by a *timer\_gettime()* call at that point in time.

#### 47255 **RETURN VALUE**

47256 If the *timer\_getoverrun()* function succeeds, it shall return the timer expiration overrun count as  
 47257 explained above.

47258 If the *timer\_gettime()* or *timer\_settime()* functions succeed, a value of 0 shall be returned.

47259 If an error occurs for any of these functions, the value -1 shall be returned, and *errno* set to  
 47260 indicate the error.

#### 47261 **ERRORS**

47262 The *timer\_getoverrun()*, *timer\_gettime()*, and *timer\_settime()* functions shall if:

47263 [EINVAL] The *timerid* argument does not correspond to an ID returned by *timer\_create()*  
 47264 but not yet deleted by *timer\_delete()*.

47265 The *timer\_settime()* function shall fail if:

47266 [EINVAL] A *value* structure specified a nanosecond value less than zero or greater than  
 47267 or equal to 1,000 million, and the *it\_value* member of that structure did not  
 47268 specify zero seconds and nanoseconds.

#### 47269 **EXAMPLES**

47270 None.

#### 47271 **APPLICATION USAGE**

47272 None.

#### 47273 **RATIONALE**

47274 The *clock\_settime()*, *timer\_settime()*, and *nanosleep()* functions are defined to truncate specified  
 47275 time values down to the resolution supported by the implementation. Values are truncated  
 47276 when set because this appears to be existing practice, and it does not seem reasonable to require  
 47277 an error in this case. Note that this is symmetric with the truncation that occurs when reading  
 47278 the time via *clock\_gettime()* or *timer\_gettime()* at a time that is not an integral multiple of the  
 47279 clock or timer resolution.

47280 This volume of IEEE Std. 1003.1-200x defines functions that allow an application to determine  
 47281 the implementation-supported resolution for the clocks and requires an implementation to  
 47282 document the resolution supported for timers and *nanosleep()* if they differ from the supported  
 47283 clock resolution. This is more of a procurement issue than a runtime application issue.

47284 **FUTURE DIRECTIONS**

47285 None.

47286 **SEE ALSO**47287 *clock\_getres()*, *timer\_create()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**time.h**>47288 **CHANGE HISTORY**

47289 First released in Issue 5. Included for alignment with the POSIX Realtime Extension.

47290 **Issue 6**47291 The *timer\_getoverrun()*, *timer\_gettime()*, and *timer\_settime()* functions are marked as part of the  
47292 Timers option.47293 The [ENOSYS] error condition has been removed as stubs need not be provided if an  
47294 implementation does not support the Timers option.47295 The [EINVAL] error condition is updated to include the following: “and the *it\_value* member of  
47296 that structure did not specify zero seconds and nanoseconds.” This change is for IEEE PASC  
47297 Interpretation 1003.1 #89.47298 The DESCRIPTION for *timer\_getoverrun()* is updated to clarify that “If no expiration signal has  
47299 been delivered for the timer, or if the Realtime Signals Extension is not supported, the return  
47300 value of *timer\_getoverrun()* is unspecified”.47301 The **restrict** keyword is added to the *timer\_settime()* prototype for alignment with the  
47302 ISO/IEC 9899:1999 standard.

47303 **NAME**

47304 times — get process and waited-for child process times

47305 **SYNOPSIS**

47306 #include &lt;sys/times.h&gt;

47307 clock\_t times(struct tms \*buffer);

47308 **DESCRIPTION**47309 The *times()* function shall fill the **tms** structure pointed to by *buffer* with time-accounting  
47310 information. The structure **tms** is defined in <sys/times.h>.

47311 All times are measured in terms of the number of clock ticks used.

47312 The times of a terminated child process are included in the *tms\_cutime* and *tms\_cstime* elements  
47313 of the parent when *wait()* or *waitpid()* returns the process ID of this terminated child. If a child  
47314 process has not waited for its children, their times shall not be included in its times.47315 • The *tms\_utime* structure member is the CPU time charged for the execution of user  
47316 instructions of the calling process.47317 • The *tms\_stime* structure member is the CPU time charged for execution by the system on  
47318 behalf of the calling process.47319 • The *tms\_cutime* structure member is the sum of the *tms\_utime* and *tms\_cutime* times of the  
47320 child processes.47321 • The *tms\_cstime* structure member is the sum of the *tms\_stime* and *tms\_cstime* times of the child  
47322 processes.47323 **RETURN VALUE**47324 Upon successful completion, *times()* shall return the elapsed real time, in clock ticks, since an  
47325 arbitrary point in the past (for example, system start-up time). This point does not change from  
47326 one invocation of *times()* within the process to another. The return value may overflow the  
47327 possible range of type **clock\_t**. If *times()* fails, **(clock\_t)-1** shall be returned and *errno* set to  
47328 indicate the error.47329 **ERRORS**

47330 No errors are defined.

47331 **EXAMPLES**47332 **Timing a Database Lookup**47333 The following example defines two functions, *start\_clock()* and *end\_clock()*, that are used to time  
47334 a lookup. It also defines variables of type **clock\_t** and **tms** to measure the duration of  
47335 transactions. The *start\_clock()* function saves the beginning times given by the *times()* function.  
47336 The *end\_clock()* function gets the ending times and prints the difference between the two times.47337 #include <sys/times.h>  
47338 #include <stdio.h>  
47339 ...  
47340 void start\_clock(void);  
47341 void end\_clock(char \*msg);  
47342 ...  
47343 static clock\_t st\_time;  
47344 static clock\_t en\_time;  
47345 static struct tms st\_cpu;  
47346 static struct tms en\_cpu;

```

47347 ...
47348 void
47349 start_clock()
47350 {
47351 st_time = times(&st_cpu);
47352 }
47353 void
47354 end_clock(char *msg)
47355 {
47356 en_time = times(&en_cpu);
47357
47358 printf(msg);
47359 printf("Real Time: %ld, User Time %ld, System Time %ld\n",
47360 en_time - st_time,
47361 en_cpu.tms_utime - st_cpu.tms_utime,
47362 en_cpu.tms_stime - st_cpu.tms_stime);
47362 }

```

#### 47363 APPLICATION USAGE

47364 Applications should use `sysconf(_SC_CLK_TCK)` to determine the number of clock ticks per  
47365 second as it may vary from system to system.

#### 47366 RATIONALE

47367 The accuracy of the times reported is intentionally left unspecified to allow implementations  
47368 flexibility in design, from uniprocessor to multiprocessor networks.

47369 The inclusion of times of child processes is recursive, so that a parent process may collect the  
47370 total times of all of its descendants. But the times of a child are only added to those of its parent  
47371 when its parent successfully waits on the child. Thus, it is not guaranteed that a parent process  
47372 can always see the total times of all its descendants; see also the discussion of the term *realtime* in  
47373 *alarm()*.

47374 If the type `clock_t` is defined to be a signed 32-bit integer, it overflows in somewhat more than a  
47375 year if there are 60 clock ticks per second, or less than a year if there are 100. There are individual  
47376 systems that run continuously for longer than that. This volume of IEEE Std. 1003.1-200x permits  
47377 an implementation to make the reference point for the returned value be the start-up time of the  
47378 process, rather than system start-up time.

47379 The term *charge* in this context has nothing to do with billing for services. The operating system  
47380 accounts for time used in this way. That information must be correct, regardless of how that  
47381 information is used.

#### 47382 FUTURE DIRECTIONS

47383 None.

#### 47384 SEE ALSO

47385 `exec`, `fork()`, `sysconf()`, `time()`, `wait()`, the Base Definitions volume of IEEE Std. 1003.1-200x,  
47386 `<sys/times.h>`

#### 47387 CHANGE HISTORY

47388 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 47389 Issue 4

47390 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 47391 • All references to the constant `CLK_TCK` are removed.

47392

- The RETURN VALUE section is updated to indicate that **(clock\_t)-1** is returned on error.

47393 **NAME**

47394            timezone — difference from UTC and local standard time

47395 **SYNOPSIS**

47396            #include <time.h>

47397 XSI        extern long timezone;

47398

47399 **DESCRIPTION**

47400            Refer to *tzset()*.

47401 **NAME**

47402 tmpfile — create a temporary file

47403 **SYNOPSIS**

47404 #include &lt;stdio.h&gt;

47405 FILE \*tmpfile(void);

47406 **DESCRIPTION**

47407 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 47408 conflict between the requirements described here and the ISO C standard is unintentional. This  
 47409 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

47410 The *tmpfile()* function shall create a temporary file and open a corresponding stream. The file is  
 47411 automatically deleted when all references to the file are closed. The file is opened as in *fopen()*  
 47412 for update (w+).

47413 CX The largest value that can be represented correctly in an object of type **off\_t** is established as the  
 47414 offset maximum in the open file description.

47415 In some implementations, a permanent file may be left behind if the process calling *tmpfile()* is  
 47416 killed while it is processing a call to *tmpfile()*.

47417 An error message may be written to standard error if the stream cannot be opened.

47418 **RETURN VALUE**

47419 Upon successful completion, *tmpfile()* shall return a pointer to the stream of the file that is  
 47420 CX created. Otherwise, it shall return a null pointer and set *errno* to indicate the error.

47421 **ERRORS**47422 The *tmpfile()* function shall fail if:47423 CX [EINTR] A signal was caught during *tmpfile()*.

47424 CX [EMFILE] {OPEN\_MAX} file descriptors are currently open in the calling process.

47425 CX [ENFILE] The maximum allowable number of files is currently open in the system.

47426 CX [ENOSPC] The directory or file system which would contain the new file cannot be  
 47427 expanded.

47428 CX [EOVERFLOW] The file is a regular file and the size of the file cannot be represented correctly  
 47429 in an object of type **off\_t**.

47430 The *tmpfile()* function may fail if:

47431 CX [EMFILE] {FOPEN\_MAX} streams are currently open in the calling process.

47432 CX [ENOMEM] Insufficient storage space is available.

47433 **EXAMPLES**47434 **Creating a Temporary File**

47435 The following example creates a temporary file for update, and returns a pointer to a stream for  
 47436 the created file in the *fp* variable.

47437 #include &lt;stdio.h&gt;

47438 ...

47439 FILE \*fp;

47440 fp = tmpfile ();

47441 **APPLICATION USAGE**

47442 It should be possible to open at least {TMP\_MAX} temporary files during the lifetime of the  
47443 program (this limit may be shared with *tmpnam()*) and there should be no limit on the number  
47444 simultaneously open other than this limit and any limit on the number of open files  
47445 ({FOPEN\_MAX}).

47446 **RATIONALE**

47447 None.

47448 **FUTURE DIRECTIONS**

47449 None.

47450 **SEE ALSO**

47451 *open()*, *tmpnam()*, *unlink()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdio.h>

47452 **CHANGE HISTORY**

47453 First released in Issue 1. Derived from Issue 1 of the SVID.

47454 **Issue 4**

47455 The argument list is explicitly defined as **void**.

47456 The [EINTR] error is moved to the “fails” part of the ERRORS section; [EMFILE], [ENFILE], and  
47457 [ENOSPC] are no longer marked as extensions; [EACCES], [ENOTDIR], and [EROFS] are  
47458 removed; and the [EMFILE] error in the “may fail” part is marked as an extension.

47459 **Issue 5**

47460 Large File Summit extensions are added.

47461 The last two paragraphs of the DESCRIPTION were included as APPLICATION USAGE notes  
47462 in previous issues.

47463 **Issue 6**

47464 Extensions beyond the ISO C standard are now marked.

47465 The following new requirements on POSIX implementations derive from alignment with the  
47466 Single UNIX Specification:

- 47467 • In the DESCRIPTION, text is added to indicate setting of the offset maximum in the open file  
47468 description. This change is to support large files.
- 47469 • In the ERRORS section, the [EOVERFLOW] condition is added. This change is to support  
47470 large files.
- 47471 • The [EMFILE] optional error condition is added.

47472 The APPLICATION USAGE section is added for alignment with the ISO/IEC 9899:1999  
47473 standard.

47474 **NAME**

47475 tmpnam — create a name for a temporary file

47476 **SYNOPSIS**

47477 #include &lt;stdio.h&gt;

47478 char \*tmpnam(char \*s);

47479 **DESCRIPTION**

47480 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 47481 conflict between the requirements described here and the ISO C standard is unintentional. This  
 47482 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

47483 The *tmpnam()* function shall generate a string that is a valid file name and that is not the same as  
 47484 the name of an existing file. The function is potentially capable of generating {TMP\_MAX}  
 47485 different strings, but any or all of them may already be in use by existing files and thus not be  
 47486 suitable return values.

47487 The *tmpnam()* function generates a different string each time it is called from the same process,  
 47488 up to {TMP\_MAX} times. If it is called more than {TMP\_MAX} times, the behavior is  
 47489 implementation-defined.

47490 The implementation shall behave as if no function defined in this volume of  
 47491 IEEE Std. 1003.1-200x calls *tmpnam()*.

47492 CX If the application uses any of the functions guaranteed to be available if either  
 47493 `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS` is defined, the application shall  
 47494 ensure that the *tmpnam()* function is called with a non-NULL parameter.

47495 **RETURN VALUE**

47496 Upon successful completion, *tmpnam()* shall return a pointer to a string. If no suitable string can  
 47497 be generated, the *tmpnam()* function shall return a null pointer.

47498 If the argument *s* is a null pointer, *tmpnam()* shall leave its result in an internal static object and  
 47499 return a pointer to that object. Subsequent calls to *tmpnam()* may modify the same object. If the  
 47500 argument *s* is not a null pointer, it is presumed to point to an array of at least {L\_tmpnam} chars;  
 47501 *tmpnam()* writes its result in that array and returns the argument as its value.

47502 **ERRORS**

47503 No errors are defined.

47504 **EXAMPLES**47505 **Generating a File Name**47506 The following example generates a unique file name and stores it in the array pointed to by *ptr*.

47507 #include &lt;stdio.h&gt;

47508 ...

47509 char filename[L\_tmpnam+1];

47510 char \*ptr;

47511 ptr = tmpnam(filename);

47512 **APPLICATION USAGE**

47513 This function only creates file names. It is the application's responsibility to create and remove  
 47514 the files.

47515 Between the time a path name is created and the file is opened, it is possible for some other  
 47516 process to create a file with the same name. Applications may find *tmpfile()* more useful.

47517 **RATIONALE**

47518 None.

47519 **FUTURE DIRECTIONS**

47520 None.

47521 **SEE ALSO**47522 *fopen()*, *open()*, *tmpnam()*, *tmpfile()*, *unlink()*, the Base Definitions volume of  
47523 IEEE Std. 1003.1-200x, <stdio.h>47524 **CHANGE HISTORY**

47525 First released in Issue 1. Derived from Issue 1 of the SVID.

47526 **Issue 5**

47527 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

47528 **Issue 6**

47529 Extensions beyond the ISO C standard are now marked.

47530 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

47531 The DESCRIPTION is expanded for alignment with the ISO/IEC 9899:1999 standard.

47532 **NAME**

47533           toascii — translate integer to a 7-bit ASCII character

47534 **SYNOPSIS**

47535 xSI       #include &lt;ctype.h&gt;

47536           int toascii(int c);

47537

47538 **DESCRIPTION**47539           The *toascii()* function shall convert its argument into a 7-bit ASCII character.47540 **RETURN VALUE**47541           The *toascii()* function shall return the value (*c* &0x7f).47542 **ERRORS**

47543           No errors are returned.

47544 **EXAMPLES**

47545           None.

47546 **APPLICATION USAGE**

47547           None.

47548 **RATIONALE**

47549           None.

47550 **FUTURE DIRECTIONS**

47551           None.

47552 **SEE ALSO**47553           *isascii()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>47554 **CHANGE HISTORY**

47555           First released in Issue 1. Derived from Issue 1 of the SVID.

47556 **NAME**

47557            tolower — transliterate uppercase characters to lowercase

47558 **SYNOPSIS**

47559            #include <ctype.h>

47560            int tolower(int c);

47561 **DESCRIPTION**

47562 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
47563 conflict between the requirements described here and the ISO C standard is unintentional. This  
47564 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

47565            The *tolower()* function has as a domain a type **int**, the value of which is representable as an  
47566 **unsigned char** or the value of EOF. If the argument has any other value, the behavior is  
47567 undefined. If the argument of *tolower()* represents an uppercase letter, and there exists a  
47568 **CX**        corresponding lowercase letter (as defined by character type information in the program locale  
47569 category *LC\_CTYPE*), the result is the corresponding lowercase letter. All other arguments in the  
47570 domain are returned unchanged.

47571 **RETURN VALUE**

47572            Upon successful completion, *tolower()* shall return the lowercase letter corresponding to the  
47573 argument passed; otherwise, it shall return the argument unchanged.

47574 **ERRORS**

47575            No errors are defined.

47576 **EXAMPLES**

47577            None.

47578 **APPLICATION USAGE**

47579            None.

47580 **RATIONALE**

47581            None.

47582 **FUTURE DIRECTIONS**

47583            None.

47584 **SEE ALSO**

47585            *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base Definitions  
47586 volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

47587 **CHANGE HISTORY**

47588            First released in Issue 1. Derived from Issue 1 of the SVID.

47589 **Issue 4**

47590            Reference to “shift information” is replaced by “character type information”.

47591            The RETURN VALUE section is added.

47592 **Issue 6**

47593            Extensions beyond the ISO C standard are now marked.

47594 **NAME**

47595           toupper — transliterate lowercase characters to uppercase

47596 **SYNOPSIS**

47597           #include &lt;ctype.h&gt;

47598           int toupper(int c);

47599 **DESCRIPTION**

47600 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
47601       conflict between the requirements described here and the ISO C standard is unintentional. This  
47602       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

47603       The *toupper()* function has as a domain a type **int**, the value of which is representable as an  
47604       **unsigned char** or the value of EOF. If the argument has any other value, the behavior is  
47605       undefined. If the argument of *toupper()* represents a lowercase letter, and there exists a  
47606 cx       corresponding uppercase letter (as defined by character type information in the program locale  
47607       category *LC\_CTYPE*), the result is the corresponding uppercase letter. All other arguments in the  
47608       domain are returned unchanged.

47609 **RETURN VALUE**

47610       Upon successful completion, *toupper()* shall return the uppercase letter corresponding to the  
47611       argument passed.

47612 **ERRORS**

47613       No errors are defined.

47614 **EXAMPLES**

47615       None.

47616 **APPLICATION USAGE**

47617       None.

47618 **RATIONALE**

47619       None.

47620 **FUTURE DIRECTIONS**

47621       None.

47622 **SEE ALSO**

47623       *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <ctype.h>, the Base Definitions  
47624       volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

47625 **CHANGE HISTORY**

47626       First released in Issue 1. Derived from Issue 1 of the SVID.

47627 **Issue 4**

47628       Reference to “shift information” is replaced by “character type information”.

47629       The RETURN VALUE section is added.

47630 **Issue 6**

47631       Extensions beyond the ISO C standard are now marked.

47632 **NAME**

47633 towctrans — character transliteration

47634 **SYNOPSIS**

47635 #include &lt;wctype.h&gt;

47636 wint\_t towctrans(wint\_t *wc*, wctrans\_t *desc*);47637 **DESCRIPTION**

47638 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
 47639 conflict between the requirements described here and the ISO C standard is unintentional. This  
 47640 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

47641 The *towctrans()* function transliterates the wide-character code *wc* using the mapping described  
 47642 by *desc*. The current setting of the *LC\_CTYPE* category should be the same as during the call to  
 47643 CX *wctrans()* that returned the value *desc*. If the value of *desc* is invalid (that is, not obtained by a  
 47644 call to *wctrans()* or *desc* is invalidated by a subsequent call to *setlocale()* that has affected  
 47645 category *LC\_CTYPE*), the result is unspecified.

47646 An application wishing to check for error situations should set *errno* to 0 before calling  
 47647 *towctrans()*. If *errno* is non-zero on return, an error has occurred.

47648 **RETURN VALUE**

47649 If successful, the *towctrans()* function shall return the mapped value of *wc* using the mapping  
 47650 described by *desc*. Otherwise, it shall return *wc* unchanged.

47651 **ERRORS**47652 The *towctrans()* function may fail if:47653 CX [EINVAL] *desc* contains an invalid transliteration descriptor.47654 **EXAMPLES**

47655 None.

47656 **APPLICATION USAGE**

47657 The strings "tolower" and "toupper" are reserved for the standard mapping names. In the  
 47658 table below, the functions in the left column are equivalent to the functions in the right column.

47659 tolower(*wc*) towctrans(*wc*, wctrans("tolower"))47660 toupper(*wc*) towctrans(*wc*, wctrans("toupper"))47661 **RATIONALE**

47662 None.

47663 **FUTURE DIRECTIONS**

47664 None.

47665 **SEE ALSO**

47666 *tolower()*, *toupper()*, *wctrans()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
 47667 <wctype.h>

47668 **CHANGE HISTORY**

47669 First released in Issue 5. Derived from ISO/IEC 9899:1990/Amendment 1:1995 (E).

47670 **Issue 6**

47671 Extensions beyond the ISO C standard are now marked.

47672 **NAME**

47673 tolower — transliterate uppercase wide-character code to lowercase

47674 **SYNOPSIS**

47675 #include <wctype.h>

47676 wint\_t tolower(wint\_t wc);

47677 **DESCRIPTION**

47678 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
47679 conflict between the requirements described here and the ISO C standard is unintentional. This  
47680 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

47681 The *tolower()* function has as a domain a type **wint\_t**, the value of which the application shall  
47682 ensure is a character representable as a **wchar\_t**, and a wide-character code corresponding to a  
47683 valid character in the current locale or the value of WEOF. If the argument has any other value,  
47684 the behavior is undefined. If the argument of *tolower()* represents an uppercase wide-character  
47685 code, and there exists a corresponding lowercase wide-character code (as defined by character  
47686 type information in the program locale category *LC\_CTYPE*), the result is the corresponding  
47687 lowercase wide-character code. All other arguments in the domain are returned unchanged.

47688 **RETURN VALUE**

47689 Upon successful completion, *tolower()* shall return the lowercase letter corresponding to the  
47690 argument passed; otherwise, it shall return the argument unchanged.

47691 **ERRORS**

47692 No errors are defined.

47693 **EXAMPLES**

47694 None.

47695 **APPLICATION USAGE**

47696 None.

47697 **RATIONALE**

47698 None.

47699 **FUTURE DIRECTIONS**

47700 None.

47701 **SEE ALSO**

47702 *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wctype.h**>, <**wchar.h**>, the  
47703 Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

47704 **CHANGE HISTORY**

47705 First released in Issue 4.

47706 **Issue 5**

47707 The following change has been made in this issue for alignment with  
47708 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 47709 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
47710 now made visible by inclusion of the header <**wctype.h**> rather than <**wchar.h**>.

47711 **Issue 6**

47712 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

47713 **NAME**

47714 towupper — transliterate lowercase wide-character code to uppercase

47715 **SYNOPSIS**

47716 #include <wctype.h>

47717 wint\_t towupper(wint\_t wc);

47718 **DESCRIPTION**

47719 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
47720 conflict between the requirements described here and the ISO C standard is unintentional. This  
47721 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

47722 The *towupper()* function has as a domain a type **wint\_t**, the value of which the application shall  
47723 ensure is a character representable as a **wchar\_t**, and a wide-character code corresponding to a  
47724 valid character in the current locale or the value of WEOF. If the argument has any other value,  
47725 the behavior is undefined. If the argument of *towupper()* represents a lowercase wide-character  
47726 code, and there exists a corresponding uppercase wide-character code (as defined by character  
47727 type information in the program locale category *LC\_CTYPE*), the result is the corresponding  
47728 uppercase wide-character code. All other arguments in the domain are returned unchanged.

47729 **RETURN VALUE**

47730 Upon successful completion, *towupper()* shall return the uppercase letter corresponding to the  
47731 argument passed. Otherwise, it shall return the argument unchanged.

47732 **ERRORS**

47733 No errors are defined.

47734 **EXAMPLES**

47735 None.

47736 **APPLICATION USAGE**

47737 None.

47738 **RATIONALE**

47739 None.

47740 **FUTURE DIRECTIONS**

47741 None.

47742 **SEE ALSO**

47743 *setlocale()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wctype.h**>, <**wchar.h**>, the  
47744 Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 7, Locale

47745 **CHANGE HISTORY**

47746 First released in Issue 4.

47747 **Issue 5**

47748 The following change has been made in this issue for alignment with  
47749 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 47750 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
47751 now made visible by inclusion of the header <**wctype.h**> rather than <**wchar.h**>.

47752 **Issue 6**

47753 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

47754 **NAME**

47755 trunc, truncf, trunc1 — round to truncated integer value

47756 **SYNOPSIS**

47757 #include <math.h>

47758 double trunc(double x);

47759 float truncf(float x);

47760 long double trunc1(long double x);

47761 **DESCRIPTION**

47762 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
47763 conflict between the requirements described here and the ISO C standard is unintentional. This  
47764 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

47765 These functions shall round their argument to the integer value, in floating format, nearest to but  
47766 no larger in magnitude than the argument.

47767 An application wishing to check for error situations should set *errno* to 0 before calling these  
47768 functions. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

47769 **RETURN VALUE**

47770 Upon successful completion, these functions shall return the truncated integer value.

47771 If *x* is  $\pm\text{Inf}$ , these functions shall return *x*.

47772 If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

47773 **ERRORS**

47774 These functions may fail if:

47775 [EDOM] The value of *x* is NaN.

47776 **EXAMPLES**

47777 None.

47778 **APPLICATION USAGE**

47779 None.

47780 **RATIONALE**

47781 None.

47782 **FUTURE DIRECTIONS**

47783 None.

47784 **SEE ALSO**

47785 The Base Definitions volume of IEEE Std. 1003.1-200x, <math.h>

47786 **CHANGE HISTORY**

47787 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

47788 **NAME**

47789 truncate — truncate a file to a specified length

47790 **SYNOPSIS**

47791 xSI #include &lt;unistd.h&gt;

47792 int truncate(const char \*path, off\_t length);

47793

47794 **DESCRIPTION**47795 The *truncate()* function shall cause the regular file named by *path* to have a size which shall be  
47796 equal to *length* bytes.47797 If the file previously was larger than *length*, the extra data is discarded. If the file was previously  
47798 shorter than *length*, its size is increased, and the extended area appears as if it were zero-filled.

47799 The application shall ensure that the process has write permission for the file.

47800 If the request would cause the file size to exceed the soft file size limit for the process, the  
47801 request shall fail and the implementation shall generate the SIGXFSZ signal for the process.47802 This function shall not modify the file offset for any open file descriptions associated with the  
47803 file. Upon successful completion, if the file size is changed, this function shall mark for update  
47804 the *st\_ctime* and *st\_mtime* fields of the file, and the S\_ISUID and S\_ISGID bits of the file mode  
47805 may be cleared.47806 **RETURN VALUE**47807 Upon successful completion, *truncate()* shall return 0. Otherwise, -1 shall be returned, and *errno*  
47808 set to indicate the error.47809 **ERRORS**47810 The *truncate()* function shall fail if:

47811 [EINTR] A signal was caught during execution.

47812 [EINVAL] The *length* argument was less than 0.

47813 [EFBIG] or [EINVAL]

47814 The *length* argument was greater than the maximum file size.

47815 [EIO] An I/O error occurred while reading from or writing to a file system.

47816 [EACCES] A component of the path prefix denies search permission, or write permission  
47817 is denied on the file.

47818 [EISDIR] The named file is a directory.

47819 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
47820 argument.

47821 [ENAMETOOLONG]

47822 The length of the *path* argument exceeds {PATH\_MAX} or a path name  
47823 component is longer than {NAME\_MAX}.47824 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.47825 [ENOTDIR] A component of the path prefix of *path* is not a directory.

47826 [EROFS] The named file resides on a read-only file system.

47827 The *truncate()* function may fail if:

47828 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
47829 resolution of the *path* argument.

47830 [ENAMETOOLONG]  
47831 Path name resolution of a symbolic link produced an intermediate result  
47832 whose length exceeds {PATH\_MAX}.

47833 **EXAMPLES**  
47834 None.

47835 **APPLICATION USAGE**  
47836 None.

47837 **RATIONALE**  
47838 None.

47839 **FUTURE DIRECTIONS**  
47840 None.

47841 **SEE ALSO**  
47842 *open()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

47843 **CHANGE HISTORY**  
47844 First released in Issue 4, Version 2.

47845 **Issue 5**  
47846 Moved from X/OPEN UNIX extension to BASE.  
47847 Large File Summit extensions are added.

47848 **Issue 6**  
47849 This reference page is split out from the *truncate()* reference page.  
47850 The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`. This  
47851 is since behavior may vary from one file system to another.  
47852 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.  
47853 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
47854 [ELOOP] error condition is added.

47855 **NAME**

47856 tsearch — search a binary search tree

47857 **SYNOPSIS**

```
47858 xSI #include <search.h>
```

```
47859 void *tsearch(const void *key, void **rootp,
47860 int (*compar)(const void *, const void *));
```

47861

47862 **DESCRIPTION**

47863 Refer to *tdelete()*.

47864 **NAME**

47865 ttyname, ttyname\_r — find path name of a terminal

47866 **SYNOPSIS**

47867 #include &lt;unistd.h&gt;

47868 char \*ttyname(int *fd*);47869 TSF int ttyname\_r(int *fd*, char \**name*, size\_t *namesize*);

47870

47871 **DESCRIPTION**

47872 The *ttyname()* function shall return a pointer to a string containing a null-terminated path name  
 47873 of the terminal associated with file descriptor *fd*. The return value may point to static data  
 47874 whose content is overwritten by each call.

47875 The *ttyname()* function need not be reentrant. A function that is not required to be reentrant is  
 47876 not required to be thread-safe.

47877 TSF The *ttyname\_r()* function stores the null-terminated path name of the terminal associated with  
 47878 the file descriptor *fd* in the character array referenced by *name*. The array is *namesize*  
 47879 characters long and should have space for the name and the terminating null character. The  
 47880 maximum length of the terminal name is {TTY\_NAME\_MAX}.

47881 **RETURN VALUE**

47882 Upon successful completion, *ttyname()* shall return a pointer to a string. Otherwise, a null  
 47883 pointer shall be returned and *errno* set to indicate the error.

47884 TSF If successful, the *ttyname\_r()* function shall return zero. Otherwise, an error number shall be  
 47885 returned to indicate the error.

47886 **ERRORS**47887 The *ttyname()* function may fail if:47888 [EBADF] The *fd* argument is not a valid file descriptor.47889 [ENOTTY] The *fd* argument does not refer to a terminal.47890 The *ttyname\_r()* function may fail if:47891 TSF [EBADF] The *fd* argument is not a valid file descriptor.47892 TSF [ENOTTY] The *fd* argument does not refer to a terminal.47893 **Notes to Reviewers**47894 *This section with side shading will not appear in the final copy. - Ed.*

47895 D1, XSH, ERN 397 points out an inconsistency for the [ERANGE] error  
 47896 condition below and suggests this be changed to [E2BIG].

47897 TSF [ERANGE] The value of *namesize* is smaller than the length of the string to be returned  
 47898 including the terminating null character.

47899 **EXAMPLES**

47900 None.

47901 **APPLICATION USAGE**

47902 None.

47903 **RATIONALE**

47904 The term *terminal* is used instead of the historical term *terminal device* in order to avoid a  
47905 reference to an undefined term.

47906 The thread-safe version places the terminal name in a user-supplied buffer and returns a non-  
47907 zero value if it fails. The non-thread-safe version may return the name in a static data area that  
47908 may be overwritten by each call.

47909 **FUTURE DIRECTIONS**

47910 None.

47911 **SEE ALSO**

47912 The Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>

47913 **CHANGE HISTORY**

47914 First released in Issue 1. Derived from Issue 1 of the SVID.

47915 **Issue 4**

47916 The <**unistd.h**> header is added to the SYNOPSIS.

47917 The statement indicating that *errno* is set on error in the RETURN VALUE section, and the errors  
47918 [EBADF] and [ENOTTY], are marked as extensions.

47919 **Issue 5**

47920 The *ttyname\_r()* function is included for alignment with the POSIX Threads Extension.

47921 A note indicating that the *ttyname()* function need not be reentrant is added to the  
47922 DESCRIPTION.

47923 **Issue 6**

47924 The *ttyname\_r()* function is marked as part of the Thread-Safe Functions option.

47925 The following new requirements on POSIX implementations derive from alignment with the  
47926 Single UNIX Specification:

- 47927 • The statement that *errno* is set on error is added.
- 47928 • The [EBADF] and [ENOTTY] optional error conditions are added.

47929 **NAME**

47930 twalk — traverse a binary search tree

47931 **SYNOPSIS**

47932 XSI #include &lt;search.h&gt;

47933 void twalk(const void \*root,  
47934 void (\*action)(const void \*, VISIT, int ));  
4793547936 **DESCRIPTION**47937 Refer to *tdelete()*.

47938 **NAME**

47939           tzname — timezone strings

47940 **SYNOPSIS**

47941           #include &lt;time.h&gt;

47942           extern char \*tzname[2];

47943 **DESCRIPTION**47944           Refer to *tzset()*.

47945 **NAME**

47946 daylight, timezone, tzname, tzset — set timezone conversion information

47947 **SYNOPSIS**

47948 #include &lt;time.h&gt;

47949 XSI extern int daylight;

47950 extern long timezone;

47951 extern char \*tzname[2];

47952 void tzset(void);

47953 **DESCRIPTION**

47954 The *tzset()* function uses the value of the environment variable *TZ* to set time conversion information used by *ctime()*, *localtime()*, *mktime()*, and *strftime()*. If *TZ* is absent from the environment, implementation-defined default timezone information is used.

47957 The *tzset()* function shall set the external variable *tzname* as follows:

47958 tzname[0] = "std";

47959 tzname[1] = "dst";

47960 where *std* and *dst* are as described in the Base Definitions volume of IEEE Std. 1003.1-200x, Chapter 8, Environment Variables.

47962 XSI The *tzset()* function also shall set the external variable *daylight* to 0 if Daylight Savings Time conversions should never be applied for the timezone in use; otherwise, non-zero. The external variable *timezone* shall be set to the difference, in seconds, between Coordinated Universal Time (UTC) and local standard time.

47966 **RETURN VALUE**47967 The *tzset()* function shall return no value.47968 **ERRORS**

47969 No errors are defined.

47970 **EXAMPLES**

47971 Example TZ variables and their timezone differences are given in the table below:

47972

47973

47974

47975

47976

47977

47978

47979

| TZ  | timezone |
|-----|----------|
| EST | 5*60*60  |
| GMT | 0*60*60  |
| JST | -9*60*60 |
| MET | -1*60*60 |
| MST | 7*60*60  |
| PST | 8*60*60  |

47980 **APPLICATION USAGE**

47981 None.

47982 **RATIONALE**

47983 None.

47984 **FUTURE DIRECTIONS**

47985 None.

47986 **SEE ALSO**

47987 *ctime()*, *localtime()*, *mktime()*, *strptime()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
47988 <**time.h**>

47989 **CHANGE HISTORY**

47990 First released in Issue 1. Derived from Issue 1 of the SVID.

47991 **Issue 4**

47992 The reference to *timezone* in the SYNOPSIS section is marked as an extension.

47993 The type of *timezone* is expanded to **extern long**.

47994 The <**time.h**> header is added to the SYNOPSIS section.

47995 The following change is incorporated for alignment with the ISO POSIX-1 standard:

47996 

- The argument list is explicitly defined as **void**.

47997 **NAME**

47998            ualarm — set the interval timer

47999 **SYNOPSIS**

48000 XSI        #include &lt;unistd.h&gt;

48001            useconds\_t ualarm(useconds\_t *useconds*, useconds\_t *interval*);

48002

48003 **DESCRIPTION**

48004        The *ualarm()* function shall cause the SIGALRM signal to be generated for the calling process  
 48005        after the number of realtime microseconds specified by the *useconds* argument has elapsed.  
 48006        When the *interval* argument is non-zero, repeated timeout notification occurs with a period in  
 48007        microseconds specified by the *interval* argument. If the notification signal, SIGALRM, is not  
 48008        caught or ignored, the calling process is terminated.

48009        Implementations may place limitations on the granularity of timer values. For each interval  
 48010        timer, if the requested timer value requires a finer granularity than the implementation supports,  
 48011        the actual timer value shall be rounded up to the next supported value.

48012        Interactions between *ualarm()* and any of the following are unspecified:

48013            *alarm()*  
 48014            *nanosleep()*  
 48015            *setitimer()*  
 48016            *timer\_create()*  
 48017            *timer\_delete()*  
 48018            *timer\_getoverrun()*  
 48019            *timer\_gettime()*  
 48020            *timer\_settime()*  
 48021            *sleep()*

48022 **RETURN VALUE**

48023        The *ualarm()* function shall return the number of microseconds remaining from the previous  
 48024        *ualarm()* call. If no timeouts are pending or if *ualarm()* has not previously been called, *ualarm()*  
 48025        shall return 0.

48026 **ERRORS**

48027        No errors are defined.

48028 **EXAMPLES**

48029        None.

48030 **APPLICATION USAGE**

48031        Applications are recommended to use *nanosleep()* if the Timers option is supported, or  
 48032        *setitimer()*, *timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, *timer\_gettime()*, or *timer\_settime()*  
 48033        instead of this function.

48034 **RATIONALE**

48035        None.

48036 **FUTURE DIRECTIONS**

48037        None.

48038 **SEE ALSO**

48039        *alarm()*, *nanosleep()*, *setitimer()*, *sleep()*, *timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, the Base  
 48040        Definitions volume of IEEE Std. 1003.1-200x, <unistd.h>

48041 **CHANGE HISTORY**

48042 First released in Issue 4, Version 2.

48043 **Issue 5**

48044 Moved from X/OPEN UNIX extension to BASE.

48045 **NAME**48046 `ulimit` — get and set process limits48047 **SYNOPSIS**48048 XSI `#include <ulimit.h>`48049 `long ulimit(int cmd, ...);`

48050

48051 **DESCRIPTION**

48052 The `ulimit()` function provides for control over process limits. The process limits that can be  
 48053 controlled by this function include the maximum size of a single file that can be written (this is  
 48054 equivalent to using `setrlimit()` with `RLIMIT_FSIZE`). The `cmd` values, defined in `<ulimit.h>`  
 48055 include:

48056 `UL_GETFSIZE` Return the file size limit (`RLIMIT_FSIZE`) of the process. The limit is in units  
 48057 of 512-byte blocks and is inherited by child processes. Files of any size can be  
 48058 read. The return value shall be the integer part of the soft file size limit  
 48059 divided by 512. If the result cannot be represented as a **long**, the result is  
 48060 unspecified.

48061 `UL_SETFSIZE` Set the file size limit for output operations of the process to the value of the  
 48062 second argument, taken as a **long**, multiplied by 512. If the result would  
 48063 overflow an `rlim_t`, the actual value set is unspecified. Any process may  
 48064 decrease its own limit, but only a process with appropriate privileges may  
 48065 increase the limit. The return value shall be the integer part of the new file size  
 48066 limit divided by 512.

48067 The `ulimit()` function shall not change the setting of `errno` if successful.

48068 As all return values are permissible in a successful situation, an application wishing to check for  
 48069 error situations should set `errno` to 0, then call `ulimit()`, and, if it returns `-1`, check to see if `errno` is  
 48070 non-zero.

48071 **RETURN VALUE**

48072 Upon successful completion, `ulimit()` shall return the value of the requested limit. Otherwise, `-1`  
 48073 shall be returned and `errno` set to indicate the error.

48074 **ERRORS**

48075 The `ulimit()` function shall fail and the limit shall be unchanged if:

48076 `[EINVAL]` The `cmd` argument is not valid.

48077 `[EPERM]` A process not having appropriate privileges attempts to increase its file size  
 48078 limit.

48079 **EXAMPLES**

48080 None.

48081 **APPLICATION USAGE**

48082 None.

48083 **RATIONALE**

48084 None.

48085 **FUTURE DIRECTIONS**

48086 None.

48087 **SEE ALSO**

48088 *getrlimit()*, *setrlimit()*, *write()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**ulimit.h**>

48089 **CHANGE HISTORY**

48090 First released in Issue 1. Derived from Issue 1 of the SVID.

48091 **Issue 4**

48092 The use of **long** is replaced by **long** in the SYNOPSIS and the DESCRIPTION sections.

48093 **Issue 4, Version 2**

48094 In the DESCRIPTION, the discussion of UL\_GETFSIZE and UL\_SETFSIZE is revised generally to distinguish between the soft and the hard file size limit of the process. For UL\_GETFSIZE, the return value is defined more precisely. For UL\_SETFSIZE, the effect on both file size limits is specified, as is the effect if the result would overflow an **rlim\_t**.

48098 **Issue 5**

48099 In the description of UL\_SETFSIZE, the text is corrected to refer to **rlim\_t** rather than the spurious **rlimit\_t**.

48101 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

48102 **NAME**

48103 umask — set and get file mode creation mask

48104 **SYNOPSIS**

48105 #include &lt;sys/stat.h&gt;

48106 mode\_t umask(mode\_t *cmask*);48107 **DESCRIPTION**

48108 The *umask()* function shall set the process' file mode creation mask to *cmask* and return the  
48109 previous value of the mask. Only the file permission bits of *cmask* (see <sys/stat.h>) are used; the  
48110 meaning of the other bits is implementation-defined.

48111 The process' file mode creation mask is used during *open()*, *creat()*, *mkdir()*, and *mkfifo()* to turn  
48112 off permission bits in the *mode* argument supplied. Bit positions that are set in *cmask* are cleared  
48113 in the mode of the created file.

48114 **RETURN VALUE**

48115 The file permission bits in the value returned by *umask()* shall be the previous value of the file  
48116 mode creation mask. The state of any other bits in that value is unspecified, except that a  
48117 subsequent call to *umask()* with the returned value as *cmask* shall leave the state of the mask the  
48118 same as its state before the first call, including any unspecified use of those bits.

48119 **ERRORS**

48120 No errors are defined.

48121 **EXAMPLES**

48122 None.

48123 **APPLICATION USAGE**

48124 None.

48125 **RATIONALE**

48126 Unsigned argument and return types for *umask()* were proposed. The return type and the  
48127 argument were both changed to **mode\_t**.

48128 Historical implementations have made use of additional bits in *cmask* for their implementation-  
48129 defined purposes. The addition of the text that the meaning of other bits of the field is  
48130 implementation-defined permits these implementations to conform to this volume of  
48131 IEEE Std. 1003.1-200x.

48132 **FUTURE DIRECTIONS**

48133 None.

48134 **SEE ALSO**

48135 *creat()*, *mkdir()*, *mkfifo()*, *open()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
48136 <sys/stat.h>, <sys/types.h>

48137 **CHANGE HISTORY**

48138 First released in Issue 1. Derived from Issue 1 of the SVID.

48139 **Issue 4**

48140 The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
48141 XSI-conformant systems.

48142 The RETURN VALUE section is expanded, in line with the ISO POSIX-1 standard, to describe  
48143 the situation with regard to additional bits in the file mode creation mask.

48144 **Issue 6**

48145 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

48146 The following new requirements on POSIX implementations derive from alignment with the  
48147 Single UNIX Specification:

- 48148 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
48149 required for conforming implementations of previous POSIX specifications, it was not  
48150 required for UNIX applications.

48151 **NAME**

48152            uname — get name of current system

48153 **SYNOPSIS**

48154            #include &lt;sys/utsname.h&gt;

48155            int uname(struct utsname \*name);

48156 **DESCRIPTION**48157            The *uname()* function shall store information identifying the current system in the structure pointed to by *name*.48159            The *uname()* function uses the **utsname** structure defined in <sys/utsname.h>.48160            The *uname()* function returns a string naming the current system in the character array *sysname*. Similarly, *nodename* contains the name that the system is known by on a communications network. The arrays *release* and *version* further identify the operating system. The array *machine* contains a name that identifies the hardware that the system is running on.

48164            The format of each member is implementation-defined.

48165 **RETURN VALUE**48166            Upon successful completion, a non-negative value shall be returned. Otherwise, -1 shall be returned and *errno* set to indicate the error.48168 **ERRORS**

48169            No errors are defined.

48170 **EXAMPLES**

48171            None.

48172 **APPLICATION USAGE**48173            The inclusion of the *nodename* member in this structure does not imply that it is sufficient information for interfacing to communications networks.48175 **RATIONALE**48176            The values of the structure members are not constrained to have any relation to the version of this volume of IEEE Std. 1003.1-200x implemented in the operating system. An application should instead depend on `_POSIX_VERSION` and related constants defined in <unistd.h>.48179            This volume of IEEE Std. 1003.1-200x does not define the sizes of the members of the structure and permits them to be of different sizes, although most implementations define them all to be the same size: eight bytes plus one byte for the string terminator. That size for *nodename* is not enough for use with many networks.48183            The *uname()* function is specific to System III, System V, and related implementations, and it does not exist in Version 7 or 4.3 BSD. The values it returns are set at system compile time in those historical implementations.48186            4.3 BSD has *gethostname()* and *gethostid()*, which return a symbolic name and a numeric value, respectively. There are related *sethostname()* and *sethostid()* functions that are used to set the values the other two functions return. The length of the host name is limited to 31 characters in most implementations and the host ID is a 32-bit integer.48190 **FUTURE DIRECTIONS**

48191            None.

48192 **SEE ALSO**

48193       The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;sys/utsname.h&gt;

48194 **CHANGE HISTORY**

48195       First released in Issue 1. Derived from Issue 1 of the SVID.

48196 **Issue 4**

48197       The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 48198       • The DESCRIPTION is changed to indicate that the format of members in the **utsname**
- 48199        structure is implementation-defined.
- 48200       • The RETURN VALUE section is updated to indicate that -1 is returned and *errno* set to
- 48201        indicate an error.

48202 **NAME**

48203 ungetc — push byte back into input stream

48204 **SYNOPSIS**

48205 #include &lt;stdio.h&gt;

48206 int ungetc(int *c*, FILE \**stream*);48207 **DESCRIPTION**

48208 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
48209 conflict between the requirements described here and the ISO C standard is unintentional. This  
48210 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

48211 The *ungetc()* function shall push the byte specified by *c* (converted to an **unsigned char**) back  
48212 onto the input stream pointed to by *stream*. The pushed-back bytes shall be returned by  
48213 subsequent reads on that stream in the reverse order of their pushing. A successful intervening  
48214 call (with the stream pointed to by *stream*) to a file-positioning function (*fseek()*, *fsetpos()*, or  
48215 *rewind()*) discards any pushed-back bytes for the stream. The external storage corresponding to  
48216 the stream is unchanged.

48217 One byte of push-back is guaranteed. If *ungetc()* is called too many times on the same stream  
48218 without an intervening read or file-positioning operation on that stream, the operation may fail.

48219 If the value of *c* equals that of the macro EOF, the operation fails and the input stream is  
48220 unchanged.

48221 A successful call to *ungetc()* shall clear the end-of-file indicator for the stream. The value of the  
48222 file-position indicator for the stream after reading or discarding all pushed-back bytes shall be  
48223 the same as it was before the bytes were pushed back. The file-position indicator is decremented  
48224 by each successful call to *ungetc()*; if its value was 0 before a call, its value is indeterminate after  
48225 the call.

48226 **RETURN VALUE**

48227 Upon successful completion, *ungetc()* shall return the byte pushed back after conversion.  
48228 Otherwise, it shall return EOF.

48229 **ERRORS**

48230 No errors are defined.

48231 **EXAMPLES**

48232 None.

48233 **APPLICATION USAGE**

48234 None.

48235 **RATIONALE**

48236 None.

48237 **FUTURE DIRECTIONS**

48238 None.

48239 **SEE ALSO**

48240 *fseek()*, *getc()*, *fsetpos()*, *read()*, *rewind()*, *setbuf()*, the Base Definitions volume of  
48241 IEEE Std. 1003.1-200x, <stdio.h>

48242 **CHANGE HISTORY**

48243 First released in Issue 1. Derived from Issue 1 of the SVID.

48244 **Issue 4**

48245 The DESCRIPTION is changed to make it clear that *ungetc()* manipulates bytes rather than  
48246 (possibly multi-byte) characters.

48247 The APPLICATION USAGE section is removed.

48248 The following changes are incorporated for alignment with the ISO C standard:

- 48249 • The *fsetpos()* function is added to the list of file-positioning functions in the DESCRIPTION.
- 48250 • Also, this issue states that the file-position indicator is decremented by each successful call to  
48251 *ungetc()*, although note that XSI-conformant systems do not distinguish between text and  
48252 binary streams. Previous issues state that the disposition of this indicator is unspecified.

48253 **NAME**

48254 ungetwc — push wide-character code back into input stream

48255 **SYNOPSIS**

48256 #include &lt;stdio.h&gt;

48257 #include &lt;wchar.h&gt;

48258 wint\_t ungetwc(wint\_t *wc*, FILE \**stream*);48259 **DESCRIPTION**

48260 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
 48261 conflict between the requirements described here and the ISO C standard is unintentional. This  
 48262 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

48263 The *ungetwc()* function shall push the character corresponding to the wide-character code  
 48264 specified by *wc* back onto the input stream pointed to by *stream*. The pushed-back characters  
 48265 shall be returned by subsequent reads on that stream in the reverse order of their pushing. A  
 48266 successful intervening call (with the stream pointed to by *stream*) to a file-positioning function  
 48267 (*fseek()*, *fsetpos()*, or *rewind()*) discards any pushed-back characters for the stream. The external  
 48268 storage corresponding to the stream is unchanged.

48269 One character of push-back is guaranteed. If *ungetwc()* is called too many times on the same  
 48270 stream without an intervening read or file-positioning operation on that stream, the operation  
 48271 may fail.

48272 If the value of *wc* equals that of the macro WEOF, the operation fails and the input stream is  
 48273 unchanged.

48274 A successful call to *ungetwc()* shall clear the end-of-file indicator for the stream. The value of the  
 48275 file-position indicator for the stream after reading or discarding all pushed-back characters shall  
 48276 be the same as it was before the characters were pushed back. The file-position indicator is  
 48277 decremented (by one or more) by each successful call to *ungetwc()*; if its value was 0 before a  
 48278 call, its value is indeterminate after the call.

48279 **RETURN VALUE**

48280 Upon successful completion, *ungetwc()* shall return the wide-character code corresponding to  
 48281 the pushed-back character. Otherwise, it shall return WEOF.

48282 **ERRORS**48283 The *ungetwc()* function may fail if:

|       |          |                                                                              |  |
|-------|----------|------------------------------------------------------------------------------|--|
| 48284 | [EILSEQ] | An invalid character sequence is detected, or a wide-character code does not |  |
| 48285 |          | correspond to a valid character.                                             |  |

48286 **EXAMPLES**

48287 None.

48288 **APPLICATION USAGE**

48289 None.

48290 **RATIONALE**

48291 None.

48292 **FUTURE DIRECTIONS**

48293 None.

48294 **SEE ALSO**

48295            *fseek()*, *fsetpos()*, *read()*, *rewind()*, *setbuf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
48296            `<stdio.h>`, `<wchar.h>`

48297 **CHANGE HISTORY**

48298            First released in Issue 4. Derived from the MSE working draft. |

48299 **Issue 5**

48300            The Optional Header (OH) marking is removed from `<stdio.h>`.

## 48301 NAME

48302 unlink — remove a directory entry

## 48303 SYNOPSIS

48304 #include &lt;unistd.h&gt;

48305 int unlink(const char \*path);

## 48306 DESCRIPTION

48307 The *unlink()* function shall remove a link to a file. If *path* names a symbolic link, *unlink()* |  
 48308 removes the symbolic link named by *path* and does not affect any file or directory named by the |  
 48309 contents of the symbolic link. Otherwise, *unlink()* removes the link named by the path name |  
 48310 pointed to by *path* and decrements the link count of the file referenced by the link.

48311 When the file's link count becomes 0 and no process has the file open, the space occupied by the |  
 48312 file shall be freed and the file shall no longer be accessible. If one or more processes have the file |  
 48313 open when the last link is removed, the link shall be removed before *unlink()* returns, but the |  
 48314 removal of the file contents shall be postponed until all references to the file are closed.

48315 The application shall ensure that the *path* argument does not name a directory unless the process |  
 48316 has appropriate privileges and the implementation supports using *unlink()* on directories.

48317 Upon successful completion, *unlink()* shall mark for update the *st\_ctime* and *st\_mtime* fields of |  
 48318 the parent directory. Also, if the file's link count is not 0, the *st\_ctime* field of the file shall be |  
 48319 marked for update.

## 48320 RETURN VALUE

48321 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* set to |  
 48322 indicate the error. If -1 is returned, the named file shall not be changed.

## 48323 ERRORS

48324 The *unlink()* function shall fail and shall not unlink the file if:

48325 [EACCES] Search permission is denied for a component of the path prefix, or write |  
 48326 permission is denied on the directory containing the directory entry to be |  
 48327 removed.

48328 [EBUSY] The file named by the *path* argument cannot be unlinked because it is being |  
 48329 used by the system or another process and the implementation considers this |  
 48330 an error.

48331 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path* |  
 48332 argument.

48333 [ENAMETOOLONG] |  
 48334 The length of the *path* argument exceeds {PATH\_MAX} or a path name |  
 48335 component is longer than {NAME\_MAX}.

48336 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

48337 [ENOTDIR] A component of the path prefix is not a directory.

48338 [EPERM] The file named by *path* is a directory, and either the calling process does not |  
 48339 have appropriate privileges, or the implementation prohibits using *unlink()* |  
 48340 on directories.

48341 XSI [EPERM] or [EACCES]

48342 The S\_ISVTX flag is set on the directory containing the file referred to by the |  
 48343 *path* argument and the caller is not the file owner, nor is the caller the |  
 48344 directory owner, nor does the caller have appropriate privileges.

48345 [EROFS] The directory entry to be unlinked is part of a read-only file system.

48346 The *unlink()* function may fail and not unlink the file if:

48347 XSI [EBUSY] The file named by *path* is a named STREAM.

48348 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
48349 resolution of the *path* argument.

48350 [ENAMETOOLONG]  
48351 As a result of encountering a symbolic link in resolution of the *path* argument,  
48352 the length of the substituted path name string exceeded {PATH\_MAX}.

48353 [ETXTBSY] The entry to be unlinked is the last directory entry to a pure procedure (shared  
48354 text) file that is being executed.

48355 **EXAMPLES**

48356 **Removing a Link to a File**

48357 The following example shows how to remove a link to a file named `/home/cnd/mod1` by  
48358 removing the entry named `/modules/pass1`.

```
48359 #include <unistd.h>
48360 char *path = "/modules/pass1";
48361 int status;
48362 ...
48363 status = unlink(path);
```

48364 **Checking for an Error**

48365 The following example fragment creates a temporary password lock file named **LOCKFILE**,  
48366 which is defined as `/etc/ptmp`, and gets a file descriptor for it. If the file cannot be opened for  
48367 writing, *unlink()* is used to remove the link between the file descriptor and **LOCKFILE**.

```
48368 #include <sys/types.h>
48369 #include <stdio.h>
48370 #include <fcntl.h>
48371 #include <errno.h>
48372 #include <unistd.h>
48373 #include <sys/stat.h>
48374 #define LOCKFILE "/etc/ptmp"
48375 int pfd; /* Integer for file descriptor returned by open call. */
48376 FILE *fpfd; /* File pointer for use in putpwent(). */
48377 ...
48378 /* Open password Lock file. If it exists, this is an error. */
48379 if ((pfd = open(LOCKFILE, O_WRONLY | O_CREAT | O_EXCL, S_IRUSR
48380 | S_IWUSR | S_IRGRP | S_IROTH)) == -1) {
48381 fprintf(stderr, "Cannot open /etc/ptmp. Try again later.\n");
48382 exit(1);
48383 }
48384 /* Lock file created, proceed with fdopen of lock file so that
48385 putpwent() can be used.
48386 */
48387 if ((fpfd = fdopen(pfd, "w")) == NULL) {
```

```

48388 close(pfd);
48389 unlink(LOCKFILE);
48390 exit(1);
48391 }

```

### 48392 **Replacing Files**

48393 The following example fragment uses *unlink()* to discard links to files, so that they can be  
48394 replaced with new versions of the files. The first call remove the link to **LOCKFILE** if an error  
48395 occurs. Successive calls remove the links to **SAVEFILE** and **PASSWDFILE** so that new links can  
48396 be created, then removes the link to **LOCKFILE** when it is no longer needed.

```

48397 #include <sys/types.h>
48398 #include <stdio.h>
48399 #include <fcntl.h>
48400 #include <errno.h>
48401 #include <unistd.h>
48402 #include <sys/stat.h>

48403 #define LOCKFILE "/etc/ptmp"
48404 #define PASSWDFILE "/etc/passwd"
48405 #define SAVEFILE "/etc/opasswd"
48406 ...
48407 /* If no change was made, assume error and leave passwd unchanged. */
48408 if (!valid_change) {
48409 fprintf(stderr, "Could not change password for user %s\n", user);
48410 unlink(LOCKFILE);
48411 exit(1);
48412 }

48413 /* Change permissions on new password file. */
48414 chmod(LOCKFILE, S_IRUSR | S_IRGRP | S_IROTH);

48415 /* Remove saved password file. */
48416 unlink(SAVEFILE);

48417 /* Save current password file. */
48418 link(PASSWDFILE, SAVEFILE);

48419 /* Remove current password file. */
48420 unlink(PASSWDFILE);

48421 /* Save new password file as current password file. */
48422 link(LOCKFILE, PASSWDFILE);

48423 /* Remove lock file. */
48424 unlink(LOCKFILE);

48425 exit(0);

```

### 48426 **APPLICATION USAGE**

48427 Applications should use *rmdir()* to remove a directory.

### 48428 **RATIONALE**

48429 Unlinking a directory is restricted to the superuser in many historical implementations for  
48430 reasons given in *link()* (see also *rename()*).

48431 The meaning of [EBUSY] in historical implementations is “mount point busy”. Since this volume  
 48432 of IEEE Std. 1003.1-200x does not cover the system administration concepts of mounting and  
 48433 unmounting, the description of the error was changed to “resource busy”. (This meaning is used  
 48434 by some device drivers when a second process tries to open an exclusive use device.) The  
 48435 wording is also intended to allow implementations to refuse to remove a directory if it is the  
 48436 root or current working directory of any process.

48437 **FUTURE DIRECTIONS**

48438 None.

48439 **SEE ALSO**

48440 *close()*, *link()*, *remove()*, *rmdir()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
 48441 <**unistd.h**>

48442 **CHANGE HISTORY**

48443 First released in Issue 1. Derived from Issue 1 of the SVID.

48444 **Issue 4**

48445 The <**unistd.h**> header is added to the SYNOPSIS section.

48446 The error [ETXTBSY] is marked as an extension.

48447 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 48448 • The type of argument *path* is changed from **char\*** to **const char\***.

48449 The following change is incorporated for alignment with the FIPS requirements:

- 48450 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
 48451 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
 48452 an extension.

48453 **Issue 4, Version 2**

48454 The entry is updated for X/OPEN UNIX conformance as follows:

- 48455 • In the DESCRIPTION, the effect is specified if *path* specifies a symbolic link.
- 48456 • In the ERRORS section, [ELOOP] is added to indicate that too many symbolic links were  
 48457 encountered during path name resolution
- 48458 • In the ERRORS section, [EPERM] or [EACCES] are added to indicate a permission check  
 48459 failure when operating on directories with S\_ISVTX set.
- 48460 • In the ERRORS section, a second [ENAMETOOLONG] condition is defined that may report  
 48461 excessive length of an intermediate result of path name resolution of a symbolic link.

48462 **Issue 5**

48463 The [EBUSY] error is added to the “may fail” part of the ERRORS section.

48464 **Issue 6**

48465 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 48466 • The [ENAMETOOLONG] error is restored as an error dependent on \_POSIX\_NO\_TRUNC.  
 48467 This is since behavior may vary from one file system to another.

48468 The following new requirements on POSIX implementations derive from alignment with the  
 48469 Single UNIX Specification:

- 48470 • In the DESCRIPTION, the effect is specified if *path* specifies a symbolic link.
- 48471 • The [ELOOP] mandatory error condition is added.

48472           • A second [ENAMETOOLONG] is added as an optional error condition.

48473           • The [ETXTBSY] optional error condition is added.

48474           The following changes were made to align with the IEEE P1003.1a draft standard:

48475           • The [ELOOP] optional error condition is added.

48476           The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48477 **NAME**

48478 unlockpt — unlock a pseudo-terminal master/slave pair

48479 **SYNOPSIS**48480 XSI `#include <stdlib.h>`48481 `int unlockpt(int fildev);`

48482

48483 **DESCRIPTION**48484 The *unlockpt()* function shall unlock the slave pseudo-terminal device associated with the  
48485 master to which *fildev* refers.48486 Portable applications shall ensure that they call *unlockpt()* before opening the slave side of a  
48487 pseudo-terminal device.48488 **RETURN VALUE**48489 Upon successful completion, *unlockpt()* shall return 0. Otherwise, it shall return -1 and set *errno*  
48490 to indicate the error.48491 **ERRORS**48492 The *unlockpt()* function may fail if:48493 [EBADF] The *fildev* argument is not a file descriptor open for writing. |48494 [EINVAL] The *fildev* argument is not associated with a master pseudo-terminal device. |48495 **EXAMPLES**

48496 None.

48497 **APPLICATION USAGE**

48498 None.

48499 **RATIONALE**

48500 None.

48501 **FUTURE DIRECTIONS**

48502 None.

48503 **SEE ALSO**48504 *grantpt()*, *open()*, *ptsname()*, the Base Definitions volume of IEEE Std. 1003.1-200x, `<stdlib.h>` |48505 **CHANGE HISTORY**

48506 First released in Issue 4, Version 2.

48507 **Issue 5**

48508 Moved from X/OPEN UNIX extension to BASE.

48509 **Issue 6**

48510 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48511 **NAME**

48512           unsetenv — remove environment variable

48513 **SYNOPSIS**

48514           #include &lt;stdlib.h&gt;

48515           int unsetenv(const char \*name);

48516 **DESCRIPTION**

48517           The *unsetenv()* function removes an environment variable from the environment of the calling  
48518           process. The *name* argument points to a string, which is the name of the variable to be removed.  
48519           The named argument shall not contain an '=' character. If the named variable does not exist in  
48520           the current environment, the environment is unchanged and the function is considered to have  
48521           completed successfully.

48522           If the application modifies *environ* or the pointers to which it points, the behavior of *unsetenv()* is  
48523           undefined. The *unsetenv()* function shall update the list of pointers to which *environ* points.

48524           The *unsetenv()* function need not be reentrant. A function that is not required to be reentrant is  
48525           not required to be thread-safe.

48526 **RETURN VALUE**

48527           Upon successful completion, zero shall be returned. Otherwise, -1 shall be returned, *errno* set to  
48528           indicate the error, and the environment shall be unchanged.

48529 **ERRORS**48530           The *unsetenv()* function shall fail if:

48531           [EINVAL]           The *name* argument is a null pointer, points to an empty string, or points to a  
48532           string containing an '=' character.

48533 **EXAMPLES**

48534           None.

48535 **APPLICATION USAGE**

48536           None.

48537 **RATIONALE**48538           Refer to the RATIONALE section in *setenv()*.48539 **FUTURE DIRECTIONS**

48540           None.

48541 **SEE ALSO**

48542           *getenv()*, *setenv()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdlib.h>, |  
48543           <sys/types.h>, <unistd.h>

48544 **CHANGE HISTORY**

48545           First released in Issue 6. Derived from the IEEE P1003.1a draft standard.

## 48546 NAME

48547            usleep — suspend execution for an interval

## 48548 SYNOPSIS

48549 XSI        #include &lt;unistd.h&gt;

48550            int usleep(useconds\_t useconds);

48551

## 48552 DESCRIPTION

48553        The *usleep()* function shall cause the calling thread to be suspended from execution until either  
 48554        the number of realtime microseconds specified by the argument *useconds* has elapsed or a signal  
 48555        is delivered to the calling thread and its action is to invoke a signal-catching function or to  
 48556        terminate the process. The suspension time may be longer than requested due to the scheduling  
 48557        of other activity by the system.

48558        The application shall ensure that the *useconds* argument is less than 1,000,000. If the value of  
 48559        *useconds* is 0, then the call has no effect.

48560        If a SIGALRM signal is generated for the calling process during execution of *usleep()* and if the  
 48561        SIGALRM signal is being ignored or blocked from delivery, it is unspecified whether *usleep()*  
 48562        returns when the SIGALRM signal is scheduled. If the signal is being blocked, it is also  
 48563        unspecified whether it remains pending after *usleep()* returns or it is discarded.

48564        If a SIGALRM signal is generated for the calling process during execution of *usleep()*, except as a  
 48565        result of a prior call to *alarm()*, and if the SIGALRM signal is not being ignored or blocked from  
 48566        delivery, it is unspecified whether that signal has any effect other than causing *usleep()* to return.

48567        If a signal-catching function interrupts *usleep()* and examines or changes either the time a  
 48568        SIGALRM is scheduled to be generated, the action associated with the SIGALRM signal, or  
 48569        whether the SIGALRM signal is blocked from delivery, the results are unspecified.

48570        If a signal-catching function interrupts *usleep()* and calls *siglongjmp()* or *longjmp()* to restore an  
 48571        environment saved prior to the *usleep()* call, the action associated with the SIGALRM signal and  
 48572        the time at which a SIGALRM signal is scheduled to be generated are unspecified. It is also  
 48573        unspecified whether the SIGALRM signal is blocked, unless the process' signal mask is restored  
 48574        as part of the environment.

48575        Implementations may place limitations on the granularity of timer values. For each interval  
 48576        timer, if the requested timer value requires a finer granularity than the implementation supports,  
 48577        the actual timer value shall be rounded up to the next supported value.

48578        Interactions between *usleep()* and any of the following are unspecified:

48579            *nanosleep()*48580            *setitimer()*48581            *timer\_create()*48582            *timer\_delete()*48583            *timer\_getoverrun()*48584            *timer\_gettime()*48585            *timer\_settime()*48586            *ualarm()*48587            *sleep()*

**48588 RETURN VALUE**

48589       Upon successful completion, *usleep()* shall return 0; otherwise, it shall return -1 and set *errno* to  
48590       indicate the error.

**48591 ERRORS**

48592       The *usleep()* function may fail if:

48593       [EINVAL]       The time interval specified 1,000,000 or more microseconds. |

**48594 EXAMPLES**

48595       None.

**48596 APPLICATION USAGE**

48597       Applications are recommended to use *nanosleep()* if the Timers option is supported, or |  
48598       *setitimer()*, *timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, *timer\_gettime()*, or *timer\_settime()*  
48599       instead of this function.

**48600 RATIONALE**

48601       None.

**48602 FUTURE DIRECTIONS**

48603       None.

**48604 SEE ALSO**

48605       *alarm()*, *getitimer()*, *nanosleep()*, *sigaction()*, *sleep()*, *timer\_create()*, *timer\_delete()*,  
48606       *timer\_getoverrun()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**> |

**48607 CHANGE HISTORY**

48608       First released in Issue 4, Version 2.

**48609 Issue 5**

48610       Moved from X/OPEN UNIX extension to BASE.

48611       The DESCRIPTION is changed to indicate that timers are now thread-based rather than  
48612       process-based.

**48613 Issue 6**

48614       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48615 **NAME**

48616 utime — set file access and modification times

48617 **SYNOPSIS**

48618 #include &lt;utime.h&gt;

48619 int utime(const char \*path, const struct utimbuf \*times);

48620 **DESCRIPTION**48621 The *utime()* function shall set the access and modification times of the file named by the *path*  
48622 argument.48623 If *times* is a null pointer, the access and modification times of the file are set to the current time.  
48624 The application shall ensure that the effective user ID of the process matches the owner of the  
48625 file, or the process has write permission to the file or has appropriate privileges, to use *utime()* in  
48626 this manner.48627 If *times* is not a null pointer, *times* is interpreted as a pointer to a **utimbuf** structure and the  
48628 access and modification times are set to the values contained in the designated structure. Only a  
48629 process with effective user ID equal to the user ID of the file or a process with appropriate  
48630 privileges may use *utime()* this way.48631 The **utimbuf** structure is defined by the header <utime.h>. The times in the structure **utimbuf**  
48632 are measured in seconds since the Epoch.48633 Upon successful completion, *utime()* shall mark the time of the last file status change, *st\_ctime*,  
48634 to be updated; see <sys/stat.h>.48635 **RETURN VALUE**48636 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* shall  
48637 be set to indicate the error, and the file times shall not be affected.48638 **ERRORS**48639 The *utime()* function shall fail if:48640 [EACCES] Search permission is denied by a component of the path prefix; or the *times* |  
48641 argument is a null pointer and the effective user ID of the process does not |  
48642 match the owner of the file, the process does not have write permission for the |  
48643 file, and the process does not have appropriate privileges. |48644 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path* |  
48645 argument. |48646 [ENAMETOOLONG] |  
48647 The length of the *path* argument exceeds {PATH\_MAX} or a path name |  
48648 component is longer than {NAME\_MAX}. |48649 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string. |

48650 [ENOTDIR] A component of the path prefix is not a directory. |

48651 [EPERM] The *times* argument is not a null pointer and the calling process' effective user |  
48652 ID does not match the owner of the file and the calling process does not have |  
48653 the appropriate privileges. |

48654 [EROFS] The file system containing the file is read-only. |

48655 The *utime()* function may fail if:48656 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during |  
48657 resolution of the *path* argument. |

48658 [ENAMETOOLONG]

48659 As a result of encountering a symbolic link in resolution of the *path* argument,  
48660 the length of the substituted path name string exceeded {PATH\_MAX}.

#### 48661 EXAMPLES

48662 None.

#### 48663 APPLICATION USAGE

48664 None.

#### 48665 RATIONALE

48666 The *actime* structure member must be present so that an application may set it, even though an  
48667 implementation may ignore it and not change the access time on the file. If an application  
48668 intends to leave one of the times of a file unchanged while changing the other, it should use  
48669 *stat()* to retrieve the file's *st\_atime* and *st\_mtime* parameters, set *actime* and *modtime* in the buffer,  
48670 and change one of them before making the *utime()* call.

#### 48671 FUTURE DIRECTIONS

48672 None.

#### 48673 SEE ALSO

48674 The Base Definitions volume of IEEE Std. 1003.1-200x, <sys/types.h>, <utime.h>

#### 48675 CHANGE HISTORY

48676 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 48677 Issue 4

48678 The <sys/types.h> header is now marked as optional (OH); this header need not be included on  
48679 XSI-conformant systems.

48680 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 48681 • The type of argument *path* is changed from **char\*** to **const char\***, and *times* is changed from  
48682 **struct utimbuf\*** to **const struct utimbuf\***.

48683 The following change is incorporated for alignment with the FIPS requirements:

- 48684 • In the ERRORS section, the condition whereby [ENAMETOOLONG] is returned if a path  
48685 name component is larger than {NAME\_MAX} is now defined as mandatory and marked as  
48686 an extension.

#### 48687 Issue 4, Version 2

48688 The ERRORS section is updated for X/OPEN UNIX conformance as follows:

- 48689 • It states that [ELOOP] is returned if too many symbolic links are encountered during path  
48690 name resolution.
- 48691 • A second [ENAMETOOLONG] condition is defined that may report excessive length of an  
48692 intermediate result of path name resolution of a symbolic link.

#### 48693 Issue 6

48694 In the SYNOPSIS, the inclusion of <sys/types.h> is no longer required.

48695 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 48696 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.  
48697 This is since behavior may vary from one file system to another.

48698 The following new requirements on POSIX implementations derive from alignment with the  
48699 Single UNIX Specification:

- 48700       • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was
- 48701       required for conforming implementations of previous POSIX specifications, it was not
- 48702       required for UNIX applications.
- 48703       • The [ELOOP] mandatory error condition is added.
- 48704       • A second [ENAMETOOLONG] is added as an optional error condition.
- 48705       The following changes were made to align with the IEEE P1003.1a draft standard:
- 48706       • The [ELOOP] optional error condition is added.
- 48707       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48708 **NAME**48709 utimes — set file access and modification times (**LEGACY**)48710 **SYNOPSIS**48711 XSI `#include <sys/time.h>`48712 `int utimes(const char *path, const struct timeval times[2]);`

48713

48714 **DESCRIPTION**

48715 The *utimes()* function shall set the access and modification times of the file pointed to by the *path*  
 48716 argument to the value of the *times* argument. The *utimes()* function allows time specifications  
 48717 accurate to the microsecond.

48718 For *utimes()*, the *times* argument is an array of **timeval** structures. The first array member  
 48719 represents the date and time of last access, and the second member represents the date and time  
 48720 of last modification. The times in the **timeval** structure are measured in seconds and  
 48721 microseconds since the Epoch, although rounding toward the nearest second may occur.

48722 If the *times* argument is a null pointer, the access and modification times of the file shall be set to  
 48723 the current time. The application shall ensure that the effective user ID of the process is the same  
 48724 as the owner of the file, or has write access to the file or appropriate privileges to use this call in  
 48725 this manner. Upon completion, *utimes()* shall mark the time of the last file status change,  
 48726 *st\_ctime*, for update.

48727 **RETURN VALUE**

48728 Upon successful completion, 0 shall be returned. Otherwise, -1 shall be returned and *errno* shall  
 48729 be set to indicate the error, and the file times shall not be affected.

48730 **ERRORS**48731 The *utimes()* function shall fail if:

48732 [EACCES] Search permission is denied by a component of the path prefix; or the *times*  
 48733 argument is a null pointer and the effective user ID of the process does not  
 48734 match the owner of the file and write access is denied.

48735 [ELOOP] A loop exists in symbolic links encountered during resolution of the *path*  
 48736 argument.

48737 [ENAMETOOLONG]  
 48738 The length of the *path* argument exceeds {PATH\_MAX} or a path name  
 48739 component is longer than {NAME\_MAX}.

48740 [ENOENT] A component of *path* does not name an existing file or *path* is an empty string.

48741 [ENOTDIR] A component of the path prefix is not a directory.

48742 [EPERM] The *times* argument is not a null pointer and the calling process' effective user  
 48743 ID has write access to the file but does not match the owner of the file and the  
 48744 calling process does not have the appropriate privileges.

48745 [EROFS] The file system containing the file is read-only.

48746 The *utimes()* function may fail if:

48747 [ELOOP] More than {SYMLOOP\_MAX} symbolic links were encountered during  
 48748 resolution of the *path* argument.

48749 [ENAMETOOLONG]  
 48750 Path name resolution of a symbolic link produced an intermediate result  
 48751 whose length exceeds {PATH\_MAX}.

48752 **EXAMPLES**

48753 None.

48754 **APPLICATION USAGE**

48755 For applications portability, the *utime()* function should be used to set file access and  
48756 modification times instead of *utimes()*.

48757 **RATIONALE**

48758 None.

48759 **FUTURE DIRECTIONS**

48760 This function may be withdrawn in a future version.

48761 **SEE ALSO**

48762 The Base Definitions volume of IEEE Std. 1003.1-200x, &lt;sys/time.h&gt;

48763 **CHANGE HISTORY**

48764 First released in Issue 4, Version 2.

48765 **Issue 5**

48766 Moved from X/OPEN UNIX extension to BASE.

48767 **Issue 6**

48768 This function is marked LEGACY.

48769 The following changes are made for alignment with the ISO POSIX-1: 1996 standard:

- 48770 • The [ENAMETOOLONG] error is restored as an error dependent on `_POSIX_NO_TRUNC`.  
48771 This is since behavior may vary from one file system to another.

48772 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

48773 The wording of the mandatory [ELOOP] error condition is updated, and a second optional  
48774 [ELOOP] error condition is added.

48775 **NAME**

48776       va\_arg, va\_copy, va\_end, va\_start — handle variable argument list |

48777 **SYNOPSIS**

48778       #include <stdarg.h>

48779       type va\_arg(va\_list ap, type);

48780       void va\_copy(va\_list dest, va\_list src); |

48781       void va\_end(va\_list ap); |

48782       void va\_start(va\_list ap, argN);

48783 **DESCRIPTION**

48784       Refer to the Base Definitions volume of IEEE Std. 1003.1-200x, <stdarg.h>. |

48785 **NAME**

48786 vfork — create new process; share virtual memory

48787 **SYNOPSIS**

48788 XSI #include &lt;unistd.h&gt;

48789 pid\_t vfork(void);

48790

48791 **DESCRIPTION**

48792 The *vfork()* function has the same effect as *fork()*, except that the behavior is undefined if the  
 48793 process created by *vfork()* either modifies any data other than a variable of type **pid\_t** used to  
 48794 store the return value from *vfork()*, or returns from the function in which *vfork()* was called, or  
 48795 calls any other function before successfully calling *\_exit()* or one of the *exec* family of functions.

48796 **RETURN VALUE**

48797 Upon successful completion, *vfork()* shall return 0 to the child process and return the process ID  
 48798 of the child process to the parent process. Otherwise, -1 shall be returned to the parent, no child  
 48799 process shall be created, and *errno* shall be set to indicate the error.

48800 **ERRORS**48801 The *vfork()* function shall fail if:

48802 [EAGAIN] The system-wide limit on the total number of processes under execution  
 48803 would be exceeded, or the system-imposed limit on the total number of  
 48804 processes under execution by a single user would be exceeded.

48805 [ENOMEM] There is insufficient swap space for the new process.

48806 **EXAMPLES**

48807 None.

48808 **APPLICATION USAGE**48809 On some systems, *vfork()* is the same as *fork()*.

48810 The *vfork()* function differs from *fork()* only in that the child process can share code and data  
 48811 with the calling process (parent process). This speeds cloning activity significantly at a risk to  
 48812 the integrity of the parent process if *vfork()* is misused.

48813 The use of *vfork()* for any purpose except as a prelude to an immediate call to a function from  
 48814 the *exec* family, or to *\_exit()*, is not advised.

48815 The *vfork()* function can be used to create new processes without fully copying the address  
 48816 space of the old process. If a forked process is simply going to call *exec*, the data space copied  
 48817 from the parent to the child by *fork()* is not used. This is particularly inefficient in a paged  
 48818 environment, making *vfork()* particularly useful. Depending upon the size of the parent's data  
 48819 space, *vfork()* can give a significant performance improvement over *fork()*.

48820 The *vfork()* function can normally be used just like *fork()*. It does not work, however, to return  
 48821 while running in the child's context from the caller of *vfork()* since the eventual return from  
 48822 *vfork()* would then return to a no longer existent stack frame. Care should be taken, also, to call  
 48823 *\_exit()* rather than *exit()* if *exec* cannot be used, since *exit()* flushes and closes standard I/O  
 48824 channels, thereby damaging the parent process' standard I/O data structures. (Even with *fork()*,  
 48825 it is wrong to call *exit()*, since buffered data would then be flushed twice.)

48826 If signal handlers are invoked in the child process after *vfork()*, they must follow the same rules  
 48827 as other code in the child process.

48828 **RATIONALE**

48829 None.

48830 **FUTURE DIRECTIONS**

48831 None.

48832 **SEE ALSO**48833 *exec*, *exit()*, *fork()*, *wait()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**unistd.h**>48834 **CHANGE HISTORY**

48835 First released in Issue 4, Version 2.

48836 **Issue 5**

48837 Moved from X/OPEN UNIX extension to BASE.

48838 **NAME**

48839 vfprintf, vprintf, vsnprintf, vsprintf — format output of a stdarg argument list

48840 **SYNOPSIS**

48841 #include &lt;stdarg.h&gt;

48842 #include &lt;stdio.h&gt;

48843 int vfprintf(FILE \*restrict stream, const char \*restrict format,  
48844 va\_list ap);

48845 int vprintf(const char \*restrict format, va\_list ap);

48846 XSI int vsnprintf(char \*restrict s, size\_t n, const char \*restrict format,  
48847 va\_list ap);

48848 int vsprintf(char \*restrict s, const char \*restrict format, va\_list ap);

48849 **DESCRIPTION**48850 CX For *vfprintf()*, *vprintf()*, and *vsprintf()*: The functionality described on this reference page is  
48851 aligned with the ISO C standard. Any conflict between the requirements described here and the  
48852 ISO C standard is unintentional. This volume of IEEE Std. 1003.1-200x defers to the ISO C  
48853 standard.48854 XSI The *vprintf()*, *vfprintf()*, *vsnprintf()*, and *vsprintf()* functions shall be the same as *printf()*,  
48855 XSI *fprintf()*, *snprintf()*, and *sprintf()* respectively, except that instead of being called with a variable  
48856 number of arguments, they are called with an argument list as defined by <stdarg.h>.48857 These functions do not invoke the *va\_end* macro. As these functions invoke the *va\_arg* macro, the  
48858 value of *ap* after the return is indeterminate.48859 **RETURN VALUE**48860 Refer to *fprintf()*.48861 **ERRORS**48862 Refer to *fprintf()*.48863 **EXAMPLES**

48864 None.

48865 **APPLICATION USAGE**48866 Applications using these functions should call *va\_end(ap)* afterwards to clean up.48867 **RATIONALE**

48868 None.

48869 **FUTURE DIRECTIONS**

48870 None.

48871 **SEE ALSO**48872 *fprintf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdarg.h>, <stdio.h>48873 **CHANGE HISTORY**

48874 First released in Issue 1. Derived from Issue 1 of the SVID.

48875 **Issue 4**

48876 The APPLICATION USAGE section is added.

48877 The FUTURE DIRECTIONS section is removed.

48878 The following changes are incorporated for alignment with the ISO C standard:

- 48879
- These functions are no longer marked as extensions.

- 48880           • The type of argument *format* is changed from **char\*** to **const char\***.
- 48881           • Reference to the **<varargs.h>** header in the DESCRIPTION is replaced by **<stdarg.h>**. The  
48882           last paragraph has also been added to indicate interactions with the *va\_arg* and *va\_end*  
48883           macros.
- 48884 **Issue 5**
- 48885           The *vsnprintf()* function is added.
- 48886 **Issue 6**
- 48887           The *vfprintf()*, *vprintf()*, *vsnprintf()*, and *vsprintf()* functions are updated for alignment with the  
48888           ISO/IEC 9899:1999 standard.

48889 **NAME**

48890 vfprintf, fprintf, vscanf — format input of a stdarg list

48891 **SYNOPSIS**

48892 #include <stdarg.h>

48893 #include <stdio.h>

48894 int vfprintf(FILE \*restrict stream, const char \*restrict format,

48895 va\_list arg);

48896 int fprintf(const char \*restrict format, va\_list arg);

48897 int vscanf(const char \*restrict s, const char \*restrict format,

48898 va\_list arg);

48899 **DESCRIPTION**

48900 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
48901 conflict between the requirements described here and the ISO C standard is unintentional. This  
48902 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

48903 These functions shall be equivalent to the *scanf()*, *fscanf()*, and *sscanf()* functions, respectively,  
48904 except that instead of being called with a variable number of arguments, they are called with an  
48905 argument list as defined by the <stdarg.h> header. These functions do not invoke the *va\_end*  
48906 macro. As these functions invoke the *va\_arg* macro, the value of *ap* after the return is  
48907 indeterminate.

48908 **RETURN VALUE**

48909 Refer to *fscanf()*.

48910 **ERRORS**

48911 Refer to *fscanf()*.

48912 **EXAMPLES**

48913 None.

48914 **APPLICATION USAGE**

48915 Applications using these functions should call *va\_end(ap)* afterwards to clean up.

48916 **RATIONALE**

48917 None.

48918 **FUTURE DIRECTIONS**

48919 None.

48920 **SEE ALSO**

48921 *fscanf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdarg.h>, <stdio.h>

48922 **CHANGE HISTORY**

48923 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

48924 **NAME**

48925 vfwprintf, vswprintf, vwprintf — wide-character formatted output of a stdarg argument list

48926 **SYNOPSIS**

48927 #include &lt;stdarg.h&gt;

48928 #include &lt;stdio.h&gt;

48929 #include &lt;wchar.h&gt;

48930 int vfwprintf(FILE \*restrict stream, const wchar\_t \*restrict format,  
48931 va\_list arg);48932 int vswprintf(wchar\_t \*restrict ws, size\_t n, const wchar\_t \*restrict format,  
48933 va\_list arg);

48934 int vwprintf(const wchar\_t \*restrict format, va\_list arg);

48935 **DESCRIPTION**48936 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
48937 conflict between the requirements described here and the ISO C standard is unintentional. This  
48938 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.48939 The *vfwprintf()*, *vswprintf()*, and *vwprintf()* functions shall be the same as *fwprintf()*, *swprintf()*,  
48940 and *wprintf()* respectively, except that instead of being called with a variable number of  
48941 arguments, they are called with an argument list as defined by <stdarg.h>.48942 These functions do not invoke the *va\_end* macro. However, as these functions do invoke the  
48943 *va\_arg* macro, the value of *ap* after the return is indeterminate.48944 **RETURN VALUE**48945 Refer to *fwprintf()*.48946 **ERRORS**48947 Refer to *fwprintf()*.48948 **EXAMPLES**

48949 None.

48950 **APPLICATION USAGE**48951 Applications using these functions should call *va\_end(ap)* afterwards to clean up.48952 **RATIONALE**

48953 None.

48954 **FUTURE DIRECTIONS**

48955 None.

48956 **SEE ALSO**48957 *fwprintf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdarg.h>, <stdio.h>,  
48958 <wchar.h>48959 **CHANGE HISTORY**48960 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
48961 (E).48962 **Issue 6**48963 The *vfwprintf()*, *vswprintf()*, and *vwprintf()* prototypes are updated for alignment with the  
48964 ISO/IEC 9899:1999 standard. ()

48965 **NAME**

48966 vfwscanf, vswscanf, vwscanf — wide-character formatted input of a stdarg list

48967 **SYNOPSIS**

48968 #include <stdarg.h>

48969 #include <stdio.h>

48970 #include <wchar.h>

48971 int vfwscanf(FILE \*restrict *stream*, const wchar\_t \*restrict *format*,  
48972 va\_list *arg*);

48973 int vswscanf(const wchar\_t \*restrict *ws*, const wchar\_t \*restrict *format*,  
48974 va\_list *arg*);

48975 int vwscanf(const wchar\_t \*restrict *format*, va\_list *arg*);

48976 **DESCRIPTION**

48977 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
48978 conflict between the requirements described here and the ISO C standard is unintentional. This  
48979 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

48980 These functions shall be equivalent to the *fwscanf()*, *swscanf()*, and *wscanf()* functions,  
48981 respectively, except that instead of being called with a variable number of arguments, they are  
48982 called with an argument list as defined by the <stdarg.h> header. These functions do not invoke  
48983 the *va\_end* macro. As these functions invoke the *va\_arg* macro, the value of *ap* after the return is  
48984 indeterminate.

48985 **RETURN VALUE**

48986 Refer to *fwscanf()*.

48987 **ERRORS**

48988 Refer to *fwscanf()*.

48989 **EXAMPLES**

48990 None.

48991 **APPLICATION USAGE**

48992 Applications using these functions should call *va\_end(ap)* afterwards to clean up.

48993 **RATIONALE**

48994 None.

48995 **FUTURE DIRECTIONS**

48996 None.

48997 **SEE ALSO**

48998 *fwscanf()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <stdarg.h>, <stdio.h>,  
48999 <wchar.h>

49000 **CHANGE HISTORY**

49001 First released in Issue 6. Derived from the ISO/IEC 9899:1999 standard.

49002 **NAME**

49003 vprintf, vsnprintf, vsprintf — format output of a stdarg argument list

49004 **SYNOPSIS**

49005 #include &lt;stdarg.h&gt;

49006 #include &lt;stdio.h&gt;

49007 int vprintf(const char \**format*, va\_list *ap*);49008 XSI int vsnprintf(char \**s*, size\_t *n*, const char \**format*, va\_list *ap*);49009 int vsprintf(char \**s*, const char \**format*, va\_list *ap*);49010 **DESCRIPTION**49011 Refer to *fprintf()*.

49012 **NAME**

49013           vscanf, vsscanf — format input of a stdarg list

49014 **SYNOPSIS**

49015           #include &lt;stdarg.h&gt;

49016           #include &lt;stdio.h&gt;

49017           int vscanf(const char \*restrict *format*, va\_list *arg*);49018           int vsscanf(const char \*restrict *s*, const char \*restrict *format*,49019                 va\_list *arg*);49020 **DESCRIPTION**49021           Refer to *vfscanf()*.

49022 **NAME**

49023 vswprintf, vwprintf — wide-character formatted output of a stdarg argument list

49024 **SYNOPSIS**

49025 #include &lt;stdarg.h&gt;

49026 #include &lt;stdio.h&gt;

49027 #include &lt;wchar.h&gt;

49028 int vswprintf(wchar\_t \*ws, size\_t n, const wchar\_t \*format,  
49029 va\_list arg);

49030 int vwprintf(const wchar\_t \*format, va\_list arg);

49031 **DESCRIPTION**49032 Refer to *vwprintf()*.

49033 **NAME**

49034           vswscanf, vwscanf — wide-character formatted input of a stdarg list

49035 **SYNOPSIS**

49036           #include &lt;stdarg.h&gt;

49037           #include &lt;stdio.h&gt;

49038           #include &lt;wchar.h&gt;

49039           int vswscanf(const wchar\_t \*restrict *ws*, const wchar\_t \*restrict *format*,  
49040                       va\_list *arg*);49041           int vwscanf(const wchar\_t \*restrict *format*, va\_list *arg*);49042 **DESCRIPTION**49043           Refer to *vfwscanf()*.

49044 **NAME**

49045 wait, waitpid — wait for a child process to stop or terminate

49046 **SYNOPSIS**

49047 #include &lt;sys/wait.h&gt;

49048 pid\_t wait(int \*stat\_loc);

49049 pid\_t waitpid(pid\_t pid, int \*stat\_loc, int options);

49050 **DESCRIPTION**

49051 The *wait()* and *waitpid()* functions allow the calling process to obtain status information  
 49052 pertaining to one of its child processes. Various options permit status information to be obtained  
 49053 for child processes that have terminated or stopped. If status information is available for two or  
 49054 more child processes, the order in which their status is reported is unspecified.

49055 The *wait()* function shall suspend execution of the calling thread until status information for one  
 49056 of the terminated child processes of the calling process is available, or until delivery of a signal  
 49057 whose action is either to execute a signal-catching function or to terminate the process. If more  
 49058 than one thread is suspended in *wait()* or *waitpid()* awaiting termination of the same process,  
 49059 exactly one thread shall return the process status at the time of the target process termination. If  
 49060 status information is available prior to the call to *wait()*, return shall be immediate.

49061 The *waitpid()* function shall behave identically to *wait()* if the *pid* argument is **(pid\_t)-1** and the  
 49062 *options* argument is 0. Otherwise, its behavior shall be modified by the values of the *pid* and  
 49063 *options* arguments.

49064 The *pid* argument specifies a set of child processes for which *status* is requested. The *waitpid()*  
 49065 function shall only return the status of a child process from this set:

- 49066 • If *pid* is equal to **(pid\_t)-1**, *status* is requested for any child process. In this respect, *waitpid()*  
 49067 is then equivalent to *wait()*.
- 49068 • If *pid* is greater than 0, it specifies the process ID of a single child process for which *status* is  
 49069 requested.
- 49070 • If *pid* is 0, *status* is requested for any child process whose process group ID is equal to that of  
 49071 the calling process.
- 49072 • If *pid* is less than **(pid\_t)-1**, *status* is requested for any child process whose process group ID  
 49073 is equal to the absolute value of *pid*.

49074 The *options* argument is constructed from the bitwise-inclusive OR of zero or more of the  
 49075 following flags, defined in the header <sys/wait.h>:

49076 XSI **WCONTINUED** The *waitpid()* function shall report the status of any continued child process  
 49077 specified by *pid* whose status has not been reported since it continued from a  
 49078 job control stop.

49079 **WNOHANG** The *waitpid()* function shall not suspend execution of the calling thread if  
 49080 *status* is not immediately available for one of the child processes specified by  
 49081 *pid*.

49082 **WUNTRACED** The status of any child processes specified by *pid* that are stopped, and whose  
 49083 status has not yet been reported since they stopped, shall also be reported to  
 49084 the requesting process.

49085 XSI If the calling process has SA\_NOCLDWAIT set or has SIGCHLD set to SIG\_IGN, and the  
 49086 process has no unwaited-for children that were transformed into zombie processes, the calling  
 49087 thread shall block until all of the children of the process containing the calling thread terminate,  
 49088 and *wait()* and *waitpid()* shall fail and set *errno* to [ECHILD].

49089 If *wait()* or *waitpid()* return because the status of a child process is available, these functions  
 49090 shall return a value equal to the process ID of the child process. In this case, if the value of the  
 49091 argument *stat\_loc* is not a null pointer, information shall be stored in the location pointed to by  
 49092 *stat\_loc*. The value stored at the location pointed to by *stat\_loc* shall be 0 if and only if the status  
 49093 returned is from a terminated child process that terminated by one of the following means:

- 49094 1. The process returned 0 from *main()*.
- 49095 2. The process called *\_exit()* or *exit()* with a *status* argument of 0.
- 49096 3. The process was terminated because the last thread in the process terminated.

49097 Regardless of its value, this information may be interpreted using the following macros, which  
 49098 are defined in `<sys/wait.h>` and evaluate to integral expressions; the *stat\_val* argument is the  
 49099 integer value pointed to by *stat\_loc*.

49100 **WIFEXITED(*stat\_val*)**

49101 Evaluates to a non-zero value if *status* was returned for a child process that terminated  
 49102 normally.

49103 **WEXITSTATUS(*stat\_val*)**

49104 If the value of **WIFEXITED(*stat\_val*)** is non-zero, this macro evaluates to the low-order 8 bits  
 49105 of the *status* argument that the child process passed to *\_exit()* or *exit()*, or the value the child  
 49106 process returned from *main()*.

49107 **WIFSIGNALED(*stat\_val*)**

49108 Evaluates to non-zero value if *status* was returned for a child process that terminated due to  
 49109 the receipt of a signal that was not caught (see `<signal.h>`).

49110 **WTERMSIG(*stat\_val*)**

49111 If the value of **WIFSIGNALED(*stat\_val*)** is non-zero, this macro evaluates to the number of  
 49112 the signal that caused the termination of the child process.

49113 **WIFSTOPPED(*stat\_val*)**

49114 Evaluates to a non-zero value if *status* was returned for a child process that is currently  
 49115 stopped.

49116 **WSTOPSIG(*stat\_val*)**

49117 If the value of **WIFSTOPPED(*stat\_val*)** is non-zero, this macro evaluates to the number of the  
 49118 signal that caused the child process to stop.

49119 XSI **WIFCONTINUED(*stat\_val*)**

49120 Evaluates to a non-zero value if *status* was returned for a child process that has continued  
 49121 from a job control stop.

49122 SPN It is unspecified whether the *status* value returned by calls to *wait()* or *waitpid()* for processes  
 49123 created by *posix\_spawn()* or *posix\_spawnnp()* may indicate a **WIFSTOPPED(*stat\_val*)** before  
 49124 subsequent calls to *wait()* or *waitpid()* indicate **WIFEXITED(*stat\_val*)** as the result of an error  
 49125 detected before the new process image starts executing.

49126 It is unspecified whether the *status* value returned by calls to *wait()* or *waitpid()* for processes  
 49127 created by *posix\_spawn()* or *posix\_spawnnp()* may indicate a **WIFSIGNALED(*stat\_val*)** if a signal is  
 49128 sent to the parent's process group after *posix\_spawn()* or *posix\_spawnnp()* is called.

49129 If the information pointed to by *stat\_loc* was stored by a call to *waitpid()* that specified the  
 49130 XSI **WUNTRACED** flag and did not specify the **WCONTINUED** flag, exactly one of the macros  
 49131 **WIFEXITED(\**stat\_loc*)**, **WIFSIGNALED(\**stat\_loc*)**, and **WIFSTOPPED(\**stat\_loc*)** shall evaluate to  
 49132 a non-zero value.

49133 If the information pointed to by *stat\_loc* was stored by a call to *waitpid()* that specified the  
 49134 XSI WUNTRACED and WCONTINUED flags, exactly one of the macros WIFEXITED(\**stat\_loc*),  
 49135 XSI WIFSIGNALED(\**stat\_loc*), WIFSTOPPED(\**stat\_loc*), and WIFCONTINUED(\**stat\_loc*) shall  
 49136 evaluate to a non-zero value.

49137 If the information pointed to by *stat\_loc* was stored by a call to *waitpid()* that did not specify the  
 49138 XSI WUNTRACED or WCONTINUED flags, or by a call to the *wait()* function, exactly one of the  
 49139 macros WIFEXITED(\**stat\_loc*) and WIFSIGNALED(\**stat\_loc*) shall evaluate to a non-zero value.

49140 If the information pointed to by *stat\_loc* was stored by a call to *waitpid()* that did not specify the  
 49141 XSI WUNTRACED flag and specified the WCONTINUED flag, or by a call to the *wait()* function,  
 49142 XSI exactly one of the macros WIFEXITED(\**stat\_loc*), WIFSIGNALED(\**stat\_loc*), and  
 49143 WIFCONTINUED(\**stat\_loc*) shall evaluate to a non-zero value.

49144 If `_POSIX_REALTIME_SIGNALS` is defined, and the implementation queues the SIGCHLD  
 49145 signal, then if *wait()* or *waitpid()* returns because the status of a child process is available, any  
 49146 pending SIGCHLD signal associated with the process ID of the child process shall be discarded.  
 49147 Any other pending SIGCHLD signals shall remain pending.

49148 Otherwise, if SIGCHLD is blocked, if *wait()* or *waitpid()* return because the status of a child  
 49149 process is available, any pending SIGCHLD signal shall be cleared unless the status of another  
 49150 child process is available.

49151 For all other conditions, it is unspecified whether child *status* will be available when a SIGCHLD  
 49152 signal is delivered.

49153 There may be additional implementation-defined circumstances under which *wait()* or *waitpid()*  
 49154 report *status*. This shall not occur unless the calling process or one of its child processes explicitly  
 49155 makes use of a non-standard extension. In these cases the interpretation of the reported *status* is  
 49156 implementation-defined.

49157 XSI If a parent process terminates without waiting for all of its child processes to terminate, the  
 49158 remaining child processes shall be assigned a new parent process ID corresponding to an  
 49159 implementation-defined system process.

#### 49160 RETURN VALUE

49161 If *wait()* or *waitpid()* returns because the status of a child process is available, these functions  
 49162 shall return a value equal to the process ID of the child process for which *status* is reported. If  
 49163 *wait()* or *waitpid()* returns due to the delivery of a signal to the calling process, `-1` shall be  
 49164 returned and *errno* set to `[EINTR]`. If *waitpid()* was invoked with `WNOHANG` set in *options*, it  
 49165 has at least one child process specified by *pid* for which *status* is not available, and *status* is not  
 49166 available for any process specified by *pid*, `0` is returned. Otherwise, `(pid_t)-1` shall be returned,  
 49167 and *errno* set to indicate the error.

#### 49168 ERRORS

49169 The *wait()* function shall fail if:

49170 `[ECHILD]` The calling process has no existing unwaited-for child processes.

49171 `[EINTR]` The function was interrupted by a signal. The value of the location pointed to  
 49172 by *stat\_loc* is undefined.

49173 The *waitpid()* function shall fail if:

49174 `[ECHILD]` The process specified by *pid* does not exist or is not a child of the calling  
 49175 process, or the process group specified by *pid* does not exist or does not have  
 49176 any member process that is a child of the calling process.

49177 [EINTR] The function was interrupted by a signal. The value of the location pointed to  
49178 by *stat\_loc* is undefined.

49179 [EINVAL] The *options* argument is not valid.

#### 49180 EXAMPLES

49181 None.

#### 49182 APPLICATION USAGE

49183 None.

#### 49184 RATIONALE

49185 A call to the *wait()* or *waitpid()* function only returns *status* on an immediate child process of the  
49186 calling process; that is, a child that was produced by a single *fork()* call (perhaps followed by an  
49187 *exec* or other function calls) from the parent. If a child produces grandchildren by further use of  
49188 *fork()*, none of those grandchildren nor any of their descendants affect the behavior of a *wait()*  
49189 from the original parent process. Nothing in this volume of IEEE Std. 1003.1-200x prevents an  
49190 implementation from providing extensions that permit a process to get *status* from a grandchild  
49191 or any other process, but a process that does not use such extensions must be guaranteed to see  
49192 *status* from only its direct children.

49193 The *waitpid()* function is provided for three reasons:

- 49194 1. To support job control
- 49195 2. To permit a non-blocking version of the *wait()* function
- 49196 3. To permit a library routine, such as *system()* or *pclose()*, to wait for its children without  
49197 interfering with other terminated children for which the process has not waited

49198 The first two of these facilities are based on the *wait3()* function provided by 4.3 BSD. The  
49199 function uses the *options* argument, which is identical to an argument to *wait3()*. The  
49200 WUNTRACED flag is used only in conjunction with job control on systems supporting job  
49201 control. Its name comes from 4.3 BSD and refers to the fact that there are two types of stopped  
49202 processes in that implementation: processes being traced via the *ptrace()* debugging facility and  
49203 (untraced) processes stopped by job control signals. Since *ptrace()* is not part of this volume of  
49204 IEEE Std. 1003.1-200x, only the second type is relevant. The name WUNTRACED was retained  
49205 because its usage is the same, even though the name is not intuitively meaningful in this context.

49206 The third reason for the *waitpid()* function is to permit independent sections of a process to  
49207 spawn and wait for children without interfering with each other. For example, the following  
49208 problem occurs in developing a portable shell, or command interpreter:

```
49209 stream = popen("/bin/true");
49210 (void) system("sleep 100");
49211 (void) pclose(stream);
```

49212 On all historical implementations, the final *pclose()* fails to reap the *wait()* *status* of the *popen()*.

49213 The status values are retrieved by macros, rather than given as specific bit encodings as they are  
49214 in most historical implementations (and thus expected by existing programs). This was  
49215 necessary to eliminate a limitation on the number of signals an implementation can support that  
49216 was inherent in the traditional encodings. This volume of IEEE Std. 1003.1-200x does require that  
49217 a *status* value of zero corresponds to a process calling *\_exit(0)*, as this is the most common  
49218 encoding expected by existing programs. Some of the macro names were adopted from 4.3 BSD.

49219 These macros syntactically operate on an arbitrary integer value. The behavior is undefined  
49220 unless that value is one stored by a successful call to *wait()* or *waitpid()* in the location pointed  
49221 to by the *stat\_loc* argument. An early proposal attempted to make this clearer by specifying each

49222 argument as *\*stat\_loc* rather than *stat\_val*. However, that did not follow the conventions of other  
 49223 specifications in this volume of IEEE Std. 1003.1-200x or traditional usage. It also could have  
 49224 implied that the argument to the macro must literally be *\*stat\_loc*; in fact, that value can be  
 49225 stored or passed as an argument to other functions before being interpreted by these macros.

49226 The extension that affects *wait()* and *waitpid()* and is common in historical implementations is  
 49227 the *ptrace()* function. It is called by a child process and causes that child to stop and return a  
 49228 *status* that appears identical to the *status* indicated by WIFSTOPPED. The *status* of *ptrace()*  
 49229 children is traditionally returned regardless of the WUNTRACED flag (or by the *wait()*  
 49230 function). Most applications do not need to concern themselves with such extensions because  
 49231 they have control over what extensions they or their children use. However, applications, such  
 49232 as command interpreters, that invoke arbitrary processes may see this behavior when those  
 49233 arbitrary processes misuse such extensions.

49234 Implementations that support *core* file creation or other implementation-defined actions on  
 49235 termination of some processes traditionally provide a bit in the *status* returned by *wait()* to  
 49236 indicate that such actions have occurred.

49237 Allowing the *wait()* family of functions to discard a pending SIGCHLD signal that is associated  
 49238 with a successfully waited-for child process puts them into the *sigwait()* and *sigwaitinfo()*  
 49239 category with respect to SIGCHLD.

49240 This definition allows implementations to treat a pending SIGCHLD signal as accepted by the  
 49241 process in *wait()*, with the same meaning of “accepted” as when that word is applied to the  
 49242 *sigwait()* family of functions.

49243 Allowing the *wait()* family of functions to behave this way permits an implementation to be able  
 49244 to deal precisely with SIGCHLD signals.

49245 In particular, an implementation that does accept (discard) the SIGCHLD signal can make the  
 49246 following guarantees regardless of the queuing depth of signals in general (the list of waitable  
 49247 children can hold the SIGCHLD queue):

- 49248 1. If a SIGCHLD signal handler is established via *sigaction()* without the SA\_RESETHAND  
 49249 flag, SIGCHLD signals can be accurately counted; that is, exactly one SIGCHLD signal will  
 49250 be delivered to or accepted by the process for every child process that terminates.
- 49251 2. A single *wait()* issued from a SIGCHLD signal handler can be guaranteed to return  
 49252 immediately with status information for a child process.
- 49253 3. When SA\_SIGINFO is requested, the SIGCHLD signal handler can be guaranteed to  
 49254 receive a non-NULL pointer to a **siginfo\_t** structure that describes a child process for  
 49255 which a wait via *waitpid()* or *waitid()* will not block or fail.
- 49256 4. The *system()* function will not cause a process's SIGCHLD handler to be called as a result of  
 49257 the *fork()/exec* executed within *system()* because *system()* will accept the SIGCHLD signal  
 49258 when it performs a *waitpid()* for its child process. This is a desirable behavior of *system()*  
 49259 so that it can be used in a library without causing side effects to the application linked with  
 49260 the library.

49261 An implementation that does not permit the *wait()* family of functions to accept (discard) a  
 49262 pending SIGCHLD signal associated with a successfully waited-for child, cannot make the  
 49263 guarantees described above for the following reasons:

#### 49264 Guarantee #1

49265 Although it might be assumed that reliable queuing of all SIGCHLD signals generated by  
 49266 the system can make this guarantee, the counter example is the case of a process that blocks  
 49267 SIGCHLD and performs an indefinite loop of *fork()/wait()* operations. If the

49268 implementation supports queued signals, then eventually the system will run out of  
 49269 memory for the queue. The guarantee cannot be made because there must be some limit to  
 49270 the depth of queuing.

49271 Guarantees #2 and #3

49272 These cannot be guaranteed unless the *wait()* family of functions accepts the SIGCHLD  
 49273 signal. Otherwise, a *fork()/wait()* executed while SIGCHLD is blocked (as in the *system()*  
 49274 function) will result in an invocation of the handler when SIGCHLD is unblocked, after the  
 49275 process has disappeared.

49276 Guarantee #4

49277 Although possible to make this guarantee, *system()* would have to set the SIGCHLD  
 49278 handler to SIG\_DFL so that the SIGCHLD signal generated by its *fork()* would be discarded  
 49279 (the SIGCHLD default action is to be ignored), then restore it to its previous setting. This  
 49280 would have the undesirable side effect of discarding all SIGCHLD signals pending to the  
 49281 process.

#### 49282 FUTURE DIRECTIONS

49283 None.

#### 49284 SEE ALSO

49285 *exec*, *exit()*, *fork()*, *waitid()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <*sys/types.h*>, |  
 49286 <*sys/wait.h*>

#### 49287 CHANGE HISTORY

49288 First released in Issue 1. Derived from Issue 1 of the SVID.

#### 49289 Issue 4

49290 The <*sys/types.h*> header is now marked as optional (OH); this header need not be included on  
 49291 XSI-conformant systems.

49292 Error return values throughout the DESCRIPTION and RETURN VALUE sections are changed  
 49293 to show the proper casting (that is, **(pid\_t)-1**).

49294 The words “If the implementation supports job control” are removed from the description of  
 49295 WUNTRACED. This is because job control is defined as mandatory for Issue 4 conforming  
 49296 implementations.

49297 The following change is incorporated for alignment with the ISO POSIX-1 standard:

- 49298 • Text describing conditions under which 0 is returned when WNOHANG is set in *options* is  
 49299 added to the RETURN VALUE section.

#### 49300 Issue 4, Version 2

49301 The *waitpid()* function is added.

49302 The following changes are incorporated in the DESCRIPTION for X/OPEN UNIX conformance:

- 49303 • The WCONTINUED *options* flag and the WIFCONTINUED(*stat\_val*) macro are added.
- 49304 • Text following the list of *options* flags explains the implications of setting the  
 49305 SA\_NOCLDWAIT signal flag, or setting SIGCHLD to SIG\_IGN.
- 49306 • Text following the list of macros, which explains what macros return non-zero values in  
 49307 certain cases, is expanded and the value of the WCONTINUED flag on the previous call to  
 49308 *waitpid()* is taken into account.

49309 **Issue 5**

49310 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

49311 **Issue 6**

49312 In the SYNOPSIS, the inclusion of `<sys/types.h>` is no longer required.

49313 The following new requirements on POSIX implementations derive from alignment with the  
49314 Single UNIX Specification:

- 49315 • The requirement to include `<sys/types.h>` has been removed. Although `<sys/types.h>` was  
49316 required for conforming implementations of previous POSIX specifications, it was not  
49317 required for UNIX applications.

49318 The following changes were made to align with the IEEE P1003.1a draft standard:

- 49319 • The processing of the SIGCHLD signal and the [ECHILD] error is clarified.

49320 The semantics of `WIFSTOPPED(stat_val)`, `WIFEXITED(stat_val)`, and `WIFSIGNALED(stat_val)`  
49321 are defined with respect to `posix_spawn()` or `posix_spawnp()` for alignment with  
49322 IEEE Std. 1003.1d-1999.

49323 The DESCRIPTION is updated for alignment with the ISO/IEC 9899:1999 standard.

49324 **NAME**

49325 waitid — wait for a child process to change state

49326 **SYNOPSIS**49327 XSI `#include <sys/wait.h>`49328 `int waitid(idtype_t idtype, id_t id, siginfo_t *infop, int options);`

49329

49330 **DESCRIPTION**

49331 The *waitid()* function shall suspend the calling thread until one child of the process containing  
 49332 the calling thread changes state. It records the current state of a child in the structure pointed to  
 49333 by *infop*. If a child process changed state prior to the call to *waitid()*, *waitid()* returns  
 49334 immediately. If more than one thread is suspended in *wait()* or *waitpid()* waiting termination of  
 49335 the same process, exactly one thread returns the process status at the time of the target process  
 49336 termination

49337 The *idtype* and *id* arguments are used to specify which children *waitid()* waits for.

49338 If *idtype* is P\_PID, *waitid()* shall wait for the child with a process ID equal to (**pid\_t**)*id*.

49339 If *idtype* is P\_PGID, *waitid()* shall wait for any child with a process group ID equal to (**pid\_t**)*id*.

49340 If *idtype* is P\_ALL, *waitid()* shall wait for any children and *id* is ignored.

49341 The *options* argument is used to specify which state changes *waitid()* shall wait for. It is formed  
 49342 by OR'ing together one or more of the following flags:

49343 WEXITED Wait for processes that have exited.

49344 WSTOPPED Status shall be returned for any child that has stopped upon receipt of a signal.

49345 WCONTINUED Status shall be returned for any child that was stopped and has been  
 49346 continued.

49347 WNOHANG Return immediately if there are no children to wait for.

49348 WNOWAIT Keep the process whose status is returned in *infop* in a waitable state. This  
 49349 shall not affect the state of the process; the process may be waited for again  
 49350 after this call completes.

49351 The application shall ensure that the *infop* argument points to a **siginfo\_t** structure. If *waitid()*  
 49352 returns because a child process was found that satisfied the conditions indicated by the  
 49353 arguments *idtype* and *options*, then the structure pointed to by *infop* shall be filled in by the  
 49354 system with the status of the process. The *si\_signo* member shall always be equal to SIGCHLD.

49355 **RETURN VALUE**49356 **Notes to Reviewers**

49357 *This section with side shading will not appear in the final copy. - Ed.*

49358 D1, XSH, ERN 416 points out an omission. The following text is proposed: "If WNOHANG was  
 49359 specified and there are no children to wait for, 0 shall be returned."

49360 If *waitid()* returns due to the change of state of one of its children, 0 shall be returned. Otherwise,  
 49361 -1 shall be returned and *errno* set to indicate the error.

49362 **ERRORS**

49363 The *waitid()* function shall fail if:

49364 [ECHILD] The calling process has no existing unwaited-for child processes.

49365 [EINTR] The *waitid()* function was interrupted by a signal. |  
49366 [EINVAL] An invalid value was specified for *options*, or *idtype* and *id* specify an invalid |  
49367 set of processes. |

49368 **EXAMPLES**

49369 None.

49370 **APPLICATION USAGE**

49371 None.

49372 **RATIONALE**

49373 None.

49374 **FUTURE DIRECTIONS**

49375 None.

49376 **SEE ALSO**

49377 *exec*, *exit()*, *wait()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <sys/wait.h> |

49378 **CHANGE HISTORY**

49379 First released in Issue 4, Version 2.

49380 **Issue 5**

49381 Moved from X/OPEN UNIX extension to BASE.

49382 The DESCRIPTION is updated for alignment with the POSIX Threads Extension.

49383 **Issue 6**

49384 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49385 **NAME**

49386       waitpid — wait for a child process to stop or terminate

49387 **SYNOPSIS**

49388       #include <sys/wait.h>

49389       pid\_t waitpid(pid\_t *pid*, int *\*stat\_loc*, int *options*);

49390 **DESCRIPTION**

49391       Refer to *wait()*.

49392 **NAME**49393 `wrtomb` — convert a wide-character code to a character (restartable)49394 **SYNOPSIS**49395 `#include <stdio.h>`49396 `size_t wrtomb(char *restrict s, wchar_t wc, mbstate_t *restrict ps);`49397 **DESCRIPTION**49398 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49399 conflict between the requirements described here and the ISO C standard is unintentional. This  
49400 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.49401 If *s* is a null pointer, the `wrtomb()` function shall be equivalent to the call:49402 `wrtomb(buf, L'\0', ps)`49403 where *buf* is an internal buffer.49404 If *s* is not a null pointer, the `wrtomb()` function shall determine the number of bytes needed to  
49405 represent the character that corresponds to the wide character given by *wc* (including any shift  
49406 sequences), and stores the resulting bytes in the array whose first element is pointed to by *s*. At  
49407 most {MB\_CUR\_MAX} bytes are stored. If *wc* is a null wide character, a null byte is stored,  
49408 preceded by any shift sequence needed to restore the initial shift state. The resulting state  
49409 described is the initial conversion state.49410 If *ps* is a null pointer, the `wrtomb()` function uses its own internal `mbstate_t` object, which is  
49411 initialized at program start-up to the initial conversion state. Otherwise, the `mbstate_t` object  
49412 pointed to by *ps* is used to completely describe the current conversion state of the associated  
49413 character sequence. The implementation shall behave as if no function defined in this volume of  
49414 IEEE Std. 1003.1-200x calls `wrtomb()`.49415 **CX** If the application uses any of the `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS`  
49416 functions, the application shall ensure that the `wrtomb()` function is called with a non-NULL *ps*  
49417 argument.49418 **XSI** The behavior of this function shall be affected by the `LC_CTYPE` category of the current locale.49419 **RETURN VALUE**49420 The `wrtomb()` function shall return the number of bytes stored in the array object (including any  
49421 shift sequences). When *wc* is not a valid wide character, an encoding error shall occur. In this  
49422 case, the function shall store the value of the macros `[EILSEQ]` in *errno* and shall return  
49423 `(size_t)-1`; the conversion state shall be undefined.49424 **ERRORS**49425 The `wrtomb()` function may fail if:49426 **CX** `[EINVAL]` *ps* points to an object that contains an invalid conversion state.49427 `[EILSEQ]` Invalid wide-character code is detected.

49428 **EXAMPLES**

49429 None.

49430 **APPLICATION USAGE**

49431 None.

49432 **RATIONALE**

49433 None.

49434 **FUTURE DIRECTIONS**

49435 None.

49436 **SEE ALSO**49437 *mbstinit()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>49438 **CHANGE HISTORY**49439 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
49440 (E).49441 **Issue 6**

49442 In the DESCRIPTION, a note on using this function in a threaded application is added.

49443 Extensions beyond the ISO C standard are now marked.

49444 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49445 The *wrtomb()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49446 **NAME**49447 `wscat` — concatenate two wide-character strings49448 **SYNOPSIS**49449 `#include <wchar.h>`49450 `wchar_t *wscat(wchar_t *restrict ws1, const wchar_t *restrict ws2);`49451 **DESCRIPTION**49452 `CX` The functionality described on this reference page is aligned with the ISO C standard. Any  
49453 conflict between the requirements described here and the ISO C standard is unintentional. This  
49454 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.49455 The `wscat()` function shall append a copy of the wide-character string pointed to by `ws2`  
49456 (including the terminating null wide-character code) to the end of the wide-character string  
49457 pointed to by `ws1`. The initial wide-character code of `ws2` overwrites the null wide-character  
49458 code at the end of `ws1`. If copying takes place between objects that overlap, the behavior is  
49459 undefined.49460 **RETURN VALUE**49461 The `wscat()` function shall return `ws1`; no return value is reserved to indicate an error.49462 **ERRORS**

49463 No errors are defined.

49464 **EXAMPLES**

49465 None.

49466 **APPLICATION USAGE**

49467 None.

49468 **RATIONALE**

49469 None.

49470 **FUTURE DIRECTIONS**

49471 None.

49472 **SEE ALSO**49473 `wscncat()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`49474 **CHANGE HISTORY**

49475 First released in Issue 4. Derived from the MSE working draft.

49476 **Issue 6**49477 The Open Group corrigenda item U040/2 has been applied. In the RETURN VALUE section, `s1`  
49478 is changed to `ws1`.49479 The `wscat()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49480 **NAME**49481 `wcschr` — wide-character string scanning operation49482 **SYNOPSIS**49483 `#include <wchar.h>`49484 `wchar_t *wcschr(const wchar_t *ws, wchar_t wc);`49485 **DESCRIPTION**

49486 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49487 conflict between the requirements described here and the ISO C standard is unintentional. This  
49488 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49489 The `wcschr()` function shall locate the first occurrence of `wc` in the wide-character string pointed  
49490 to by `ws`. The application shall ensure that the value of `wc` is a character representable as a type  
49491 **wchar\_t** and a wide-character code corresponding to a valid character in the current locale. The  
49492 terminating null wide-character code is considered to be part of the wide-character string.

49493 **RETURN VALUE**

49494 Upon completion, `wcschr()` shall return a pointer to the wide-character code, or a null pointer if  
49495 the wide-character code is not found.

49496 **ERRORS**

49497 No errors are defined.

49498 **EXAMPLES**

49499 None.

49500 **APPLICATION USAGE**

49501 None.

49502 **RATIONALE**

49503 None.

49504 **FUTURE DIRECTIONS**

49505 None.

49506 **SEE ALSO**49507 `wcsrchr()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`49508 **CHANGE HISTORY**

49509 First released in Issue 4. Derived from the MSE working draft.

49510 **Issue 6**

49511 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49512 **NAME**

49513           wcsncmp — compare two wide-character strings

49514 **SYNOPSIS**

49515           #include &lt;wchar.h&gt;

49516           int wcsncmp(const wchar\_t \*ws1, const wchar\_t \*ws2);

49517 **DESCRIPTION**

49518 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
49519 conflict between the requirements described here and the ISO C standard is unintentional. This  
49520 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49521       The *wcsncmp()* function shall compare the wide-character string pointed to by *ws1* to the wide-  
49522 character string pointed to by *ws2*.

49523       The sign of a non-zero return value is determined by the sign of the difference between the  
49524 values of the first pair of wide-character codes that differ in the objects being compared.

49525 **RETURN VALUE**

49526       Upon completion, *wcsncmp()* shall return an integer greater than, equal to, or less than 0, if the  
49527 wide-character string pointed to by *ws1* is greater than, equal to, or less than the wide-character  
49528 string pointed to by *ws2*, respectively.

49529 **ERRORS**

49530       No errors are defined.

49531 **EXAMPLES**

49532       None.

49533 **APPLICATION USAGE**

49534       None.

49535 **RATIONALE**

49536       None.

49537 **FUTURE DIRECTIONS**

49538       None.

49539 **SEE ALSO**49540       *wcsncmp()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <wchar.h>49541 **CHANGE HISTORY**

49542       First released in Issue 4. Derived from the MSE working draft.

49543 **NAME**

49544 wscoll — wide-character string comparison using collating information

49545 **SYNOPSIS**

49546 #include &lt;wchar.h&gt;

49547 int wscoll(const wchar\_t \*ws1, const wchar\_t \*ws2);

49548 **DESCRIPTION**49549 CX The functionality described on this reference page is aligned with the ISO C standard. Any  
49550 conflict between the requirements described here and the ISO C standard is unintentional. This  
49551 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.49552 The *wscoll()* function shall compare the wide-character string pointed to by *ws1* to the wide-  
49553 character string pointed to by *ws2*, both interpreted as appropriate to the *LC\_COLLATE* category  
49554 of the current locale.49555 CX The *wscoll()* function shall not change the setting of *errno* if successful.49556 An application wishing to check for error situations should set *errno* to 0 before calling *wscoll()*.  
49557 If *errno* is non-zero on return, an error has occurred.49558 **RETURN VALUE**49559 Upon successful completion, *wscoll()* shall return an integer greater than, equal to, or less than  
49560 0, according to whether the wide-character string pointed to by *ws1* is greater than, equal to, or  
49561 less than the wide-character string pointed to by *ws2*, when both are interpreted as appropriate  
49562 CX to the current locale. On error, *wscoll()* may set *errno*, but no return value is reserved to indicate  
49563 an error.49564 **ERRORS**49565 The *wscoll()* function may fail if:49566 [EINVAL] The *ws1* or *ws2* arguments contain wide-character codes outside the domain of  
49567 the collating sequence.49568 **EXAMPLES**

49569 None.

49570 **APPLICATION USAGE**49571 The *wcsxfrm()* and *wscmp()* functions should be used for sorting large lists.49572 **RATIONALE**

49573 None.

49574 **FUTURE DIRECTIONS**

49575 None.

49576 **SEE ALSO**49577 *wscmp()*, *wcsxfrm()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <wchar.h>49578 **CHANGE HISTORY**

49579 First released in Issue 4. Derived from the MSE working draft.

49580 **Issue 5**

49581 Moved from ENHANCED I18N to BASE and the [ENOSYS] error is removed.

49582 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

49583 **NAME**

49584           wcscpy — copy a wide-character string

49585 **SYNOPSIS**

49586           #include &lt;wchar.h&gt;

49587           wchar\_t \*wcscpy(wchar\_t \*restrict *ws1*, const wchar\_t \*restrict *ws2*);49588 **DESCRIPTION**

49589 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
49590 conflict between the requirements described here and the ISO C standard is unintentional. This  
49591 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49592       The *wcscpy()* function shall copy the wide-character string pointed to by *ws2* (including the  
49593 terminating null wide-character code) into the array pointed to by *ws1*. If copying takes place  
49594 between objects that overlap, the behavior is undefined.

49595 **RETURN VALUE**49596       The *wcscpy()* function shall return *ws1*; no return value is reserved to indicate an error.49597 **ERRORS**

49598       No errors are defined.

49599 **EXAMPLES**

49600       None.

49601 **APPLICATION USAGE**

49602       None.

49603 **RATIONALE**

49604       None.

49605 **FUTURE DIRECTIONS**

49606       None.

49607 **SEE ALSO**49608       *wscncpy()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>49609 **CHANGE HISTORY**

49610       First released in Issue 4. Derived from the MSE working draft.

49611 **Issue 6**49612       The *wcscpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49613 **NAME**

49614 wscspn — get length of a complementary wide substring

49615 **SYNOPSIS**

49616 #include &lt;wchar.h&gt;

49617 size\_t wscspn(const wchar\_t \*ws1, const wchar\_t \*ws2);

49618 **DESCRIPTION**

49619 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49620 conflict between the requirements described here and the ISO C standard is unintentional. This  
49621 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49622 The *wscspn()* function shall compute the length of the maximum initial segment of the wide-  
49623 character string pointed to by *ws1* which consists entirely of wide-character codes *not* from the  
49624 wide-character string pointed to by *ws2*.

49625 **RETURN VALUE**

49626 The *wscspn()* function shall return the length of the initial substring of *ws1*; no return value is  
49627 reserved to indicate an error.

49628 **ERRORS**

49629 No errors are defined.

49630 **EXAMPLES**

49631 None.

49632 **APPLICATION USAGE**

49633 None.

49634 **RATIONALE**

49635 None.

49636 **FUTURE DIRECTIONS**

49637 None.

49638 **SEE ALSO**49639 *wcspn()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <wchar.h>49640 **CHANGE HISTORY**

49641 First released in Issue 4. Derived from the MSE working draft.

49642 **Issue 5**

49643 The RETURN VALUE section is updated to indicate that *wscspn()* returns the length of *ws1*,  
49644 rather than *ws1* itself.

49645 **NAME**49646 `wcsftime` — convert date and time to a wide-character string49647 **SYNOPSIS**49648 `#include <wchar.h>`49649 `size_t wcsftime(wchar_t *restrict wcs, size_t maxsize,`  
49650 `const wchar_t *restrict format, const struct tm *restrict timptr);`49651 **DESCRIPTION**49652 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49653 conflict between the requirements described here and the ISO C standard is unintentional. This  
49654 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.49655 The `wcsftime()` function shall be equivalent to the `strftime()` function, except that:

- 49656
- The argument `wcs` points to the initial element of an array of wide characters into which the  
49657 generated output is to be placed.
  - The argument `maxsize` indicates the maximum number of wide characters to be placed in the  
49658 output array.
  - The argument `format` is a wide-character string and the conversion specifications are replaced  
49661 by corresponding sequences of wide characters.
  - The return value indicates the number of wide characters placed in the output array.

49662 If copying takes place between objects that overlap, the behavior is undefined.

49664 **RETURN VALUE**49665 If the total number of resulting wide-character codes including the terminating null wide-  
49666 character code is no more than `maxsize`, `wcsftime()` shall return the number of wide-character  
49667 codes placed into the array pointed to by `wcs`, not including the terminating null wide-character  
49668 code.49669 **ERRORS**

49670 No errors are defined.

49671 **EXAMPLES**

49672 None.

49673 **APPLICATION USAGE**

49674 None.

49675 **RATIONALE**

49676 None.

49677 **FUTURE DIRECTIONS**

49678 None.

49679 **SEE ALSO**49680 `strftime()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`49681 **CHANGE HISTORY**

49682 First released in Issue 4.

49683 **Issue 5**

49684 Moved from ENHANCED I18N to BASE and the [ENOSYS] error is removed.

49685 Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, the type of the `format`  
49686 argument is changed from `const char*` to `const wchar_t*`.

49687 **Issue 6**

49688

The *wcsftime()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49689 **NAME**49690 `wcslen` — get wide-character string length49691 **SYNOPSIS**49692 `#include <wchar.h>`49693 `size_t wcslen(const wchar_t *ws);`49694 **DESCRIPTION**

49695 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
49696 conflict between the requirements described here and the ISO C standard is unintentional. This  
49697 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49698 The `wcslen()` function shall compute the number of wide-character codes in the wide-character  
49699 string to which *ws* points, not including the terminating null wide-character code.

49700 **RETURN VALUE**

49701 The `wcslen()` function shall return the length of *ws*; no return value is reserved to indicate an  
49702 error.

49703 **ERRORS**

49704 No errors are defined.

49705 **EXAMPLES**

49706 None.

49707 **APPLICATION USAGE**

49708 None.

49709 **RATIONALE**

49710 None.

49711 **FUTURE DIRECTIONS**

49712 None.

49713 **SEE ALSO**49714 The Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`49715 **CHANGE HISTORY**

49716 First released in Issue 4. Derived from the MSE working draft.

49717 **NAME**

49718        `wcsncat` — concatenate a wide-character string with part of another

49719 **SYNOPSIS**

49720        `#include <wchar.h>`

49721        `wchar_t *wcsncat(wchar_t *restrict ws1, const wchar_t *restrict ws2,`  
49722            `size_t n);`

49723 **DESCRIPTION**

49724 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
49725        conflict between the requirements described here and the ISO C standard is unintentional. This  
49726        volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49727        The `wcsncat()` function shall append not more than *n* wide-character codes (a null wide-  
49728        character code and wide-character codes that follow it are not appended) from the array pointed  
49729        to by *ws2* to the end of the wide-character string pointed to by *ws1*. The initial wide-character  
49730        code of *ws2* overwrites the null wide-character code at the end of *ws1*. A terminating null wide-  
49731        character code is always appended to the result. If copying takes place between objects that  
49732        overlap, the behavior is undefined.

49733 **RETURN VALUE**

49734        The `wcsncat()` function shall return *ws1*; no return value is reserved to indicate an error.

49735 **ERRORS**

49736        No errors are defined.

49737 **EXAMPLES**

49738        None.

49739 **APPLICATION USAGE**

49740        None.

49741 **RATIONALE**

49742        None.

49743 **FUTURE DIRECTIONS**

49744        None.

49745 **SEE ALSO**

49746        `wscat()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`

49747 **CHANGE HISTORY**

49748        First released in Issue 4. Derived from the MSE working draft.

49749 **Issue 6**

49750        The `wcsncat()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49751 **NAME**49752 `wcsncmp` — compare part of two wide-character strings49753 **SYNOPSIS**49754 `#include <wchar.h>`49755 `int wcsncmp(const wchar_t *ws1, const wchar_t *ws2, size_t n);`49756 **DESCRIPTION**

49757 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49758 conflict between the requirements described here and the ISO C standard is unintentional. This  
49759 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49760 The `wcsncmp()` function shall compare not more than *n* wide-character codes (wide-character  
49761 codes that follow a null wide-character code are not compared) from the array pointed to by *ws1*  
49762 to the array pointed to by *ws2*.

49763 The sign of a non-zero return value is determined by the sign of the difference between the  
49764 values of the first pair of wide-character codes that differ in the objects being compared.

49765 **RETURN VALUE**

49766 Upon successful completion, `wcsncmp()` shall return an integer greater than, equal to, or less  
49767 than 0, if the possibly null-terminated array pointed to by *ws1* is greater than, equal to, or less  
49768 than the possibly null-terminated array pointed to by *ws2*, respectively.

49769 **ERRORS**

49770 No errors are defined.

49771 **EXAMPLES**

49772 None.

49773 **APPLICATION USAGE**

49774 None.

49775 **RATIONALE**

49776 None.

49777 **FUTURE DIRECTIONS**

49778 None.

49779 **SEE ALSO**49780 `wscmp()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`49781 **CHANGE HISTORY**

49782 First released in Issue 4. Derived from the MSE working draft.

49783 **NAME**49784 `wcsncpy` — copy part of a wide-character string49785 **SYNOPSIS**49786 `#include <wchar.h>`49787 `wchar_t *wcsncpy(wchar_t *restrict ws1, const wchar_t *restrict ws2,`  
49788 `size_t n);`49789 **DESCRIPTION**49790 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49791 conflict between the requirements described here and the ISO C standard is unintentional. This  
49792 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.49793 The `wcsncpy()` function shall copy not more than *n* wide-character codes (wide-character codes  
49794 that follow a null wide-character code are not copied) from the array pointed to by *ws2* to the  
49795 array pointed to by *ws1*. If copying takes place between objects that overlap, the behavior is  
49796 undefined.49797 If the array pointed to by *ws2* is a wide-character string that is shorter than *n* wide-character  
49798 codes, null wide-character codes are appended to the copy in the array pointed to by *ws1*, until *n*  
49799 wide-character codes in all are written.49800 **RETURN VALUE**49801 The `wcsncpy()` function shall return *ws1*; no return value is reserved to indicate an error.49802 **ERRORS**

49803 No errors are defined.

49804 **EXAMPLES**

49805 None.

49806 **APPLICATION USAGE**49807 If there is no null wide-character code in the first *n* wide-character codes of the array pointed to  
49808 by *ws2*, the result is not null-terminated.49809 **RATIONALE**

49810 None.

49811 **FUTURE DIRECTIONS**

49812 None.

49813 **SEE ALSO**49814 `wscpy()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`49815 **CHANGE HISTORY**

49816 First released in Issue 4. Derived from the MSE working draft.

49817 **Issue 6**49818 The `wcsncpy()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49819 **NAME**

49820 `wcpbrk` — scan wide-character string for a wide-character code

49821 **SYNOPSIS**

49822 `#include <wchar.h>`

49823 `wchar_t *wcpbrk(const wchar_t *ws1, const wchar_t *ws2);`

49824 **DESCRIPTION**

49825 *cx* The functionality described on this reference page is aligned with the ISO C standard. Any  
49826 conflict between the requirements described here and the ISO C standard is unintentional. This  
49827 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49828 The `wcpbrk()` function shall locate the first occurrence in the wide-character string pointed to by  
49829 *ws1* of any wide-character code from the wide-character string pointed to by *ws2*.

49830 **RETURN VALUE**

49831 Upon successful completion, `wcpbrk()` shall return a pointer to the wide-character code or a null  
49832 pointer if no wide-character code from *ws2* occurs in *ws1*.

49833 **ERRORS**

49834 No errors are defined.

49835 **EXAMPLES**

49836 None.

49837 **APPLICATION USAGE**

49838 None.

49839 **RATIONALE**

49840 None.

49841 **FUTURE DIRECTIONS**

49842 None.

49843 **SEE ALSO**

49844 `wchr()`, `wchr()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`

49845 **CHANGE HISTORY**

49846 First released in Issue 4. Derived from the MSE working draft.

49847 **NAME**

49848       wcsrchr — wide-character string scanning operation

49849 **SYNOPSIS**

49850       #include &lt;wchar.h&gt;

49851       wchar\_t \*wcsrchr(const wchar\_t \*ws, wchar\_t wc);

49852 **DESCRIPTION**

49853 cx       The functionality described on this reference page is aligned with the ISO C standard. Any  
49854       conflict between the requirements described here and the ISO C standard is unintentional. This  
49855       volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49856       The *wcsrchr()* function shall locate the last occurrence of *wc* in the wide-character string pointed  
49857       to by *ws*. The application shall ensure that the value of *wc* is a character representable as a type  
49858       **wchar\_t** and a wide-character code corresponding to a valid character in the current locale. The  
49859       terminating null wide-character code is considered to be part of the wide-character string.

49860 **RETURN VALUE**

49861       Upon successful completion, *wcsrchr()* shall return a pointer to the wide-character code or a null  
49862       pointer if *wc* does not occur in the wide-character string.

49863 **ERRORS**

49864       No errors are defined.

49865 **EXAMPLES**

49866       None.

49867 **APPLICATION USAGE**

49868       None.

49869 **RATIONALE**

49870       None.

49871 **FUTURE DIRECTIONS**

49872       None.

49873 **SEE ALSO**49874       *wchr()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>49875 **CHANGE HISTORY**

49876       First released in Issue 4. Derived from the MSE working draft.

49877 **Issue 6**

49878       The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49879 **NAME**49880 `wcsrtombs` — convert a wide-character string to a character string (restartable)49881 **SYNOPSIS**49882 `#include <wchar.h>`49883 `size_t wcsrtombs(char *restrict dst, const wchar_t **restrict src,`  
49884 `size_t len, mbstate_t *restrict ps);`49885 **DESCRIPTION**49886 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49887 conflict between the requirements described here and the ISO C standard is unintentional. This  
49888 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.49889 The `wcsrtombs()` function shall convert a sequence of wide characters from the array indirectly  
49890 pointed to by `src` into a sequence of corresponding characters, beginning in the conversion state  
49891 described by the object pointed to by `ps`. If `dst` is not a null pointer, the converted characters are  
49892 then stored into the array pointed to by `dst`. Conversion continues up to and including a  
49893 terminating null wide character, which is also stored. Conversion stops earlier in the following  
49894 cases:

- 49895
- When a code is reached that does not correspond to a valid character
  - When the next character would exceed the limit of `len` total bytes to be stored in the array  
49896 pointed to by `dst` (and `dst` is not a null pointer)
- 49897

49898 Each conversion takes place as if by a call to the `wcrtomb()` function.49899 If `dst` is not a null pointer, the pointer object pointed to by `src` is assigned either a null pointer (if  
49900 conversion stopped due to reaching a terminating null wide character) or the address just past  
49901 the last wide character converted (if any). If conversion stopped due to reaching a terminating  
49902 null wide character, the resulting state described is the initial conversion state.49903 If `ps` is a null pointer, the `wcsrtombs()` function uses its own internal `mbstate_t` object, which is  
49904 initialized at program start-up to the initial conversion state. Otherwise, the `mbstate_t` object  
49905 pointed to by `ps` is used to completely describe the current conversion state of the associated  
49906 character sequence. The implementation shall behave as if no function defined in this volume of  
49907 IEEE Std. 1003.1-200x calls `wcsrtombs()`.49908 **CX** If the application uses any of the `_POSIX_THREAD_SAFE_FUNCTIONS` or `_POSIX_THREADS`  
49909 functions, the application shall ensure that the `wcsrtombs()` function is called with a non-NULL  
49910 `ps` argument.49911 **XSI** The behavior of this function shall be affected by the `LC_CTYPE` category of the current locale.49912 **RETURN VALUE**49913 If conversion stops because a code is reached that does not correspond to a valid character, an  
49914 encoding error occurs. In this case, the `wcsrtombs()` function shall store the value of the macro  
49915 `[EILSEQ]` in `errno` and return `(size_t)-1`; the conversion state is undefined. Otherwise, it shall  
49916 return the number of bytes in the resulting character sequence, not including the terminating  
49917 null (if any).49918 **ERRORS**49919 The `wcsrtombs()` function may fail if:

- 49920
- CX**
- `[EINVAL]`
- `ps`
- points to an object that contains an invalid conversion state.
- 
- 49921
- `[EILSEQ]`
- A wide-character code does not correspond to a valid character.

49922 **EXAMPLES**

49923 None.

49924 **APPLICATION USAGE**

49925 None.

49926 **RATIONALE**

49927 None.

49928 **FUTURE DIRECTIONS**

49929 None.

49930 **SEE ALSO**49931 *mbsinit()*, *wcrtomb()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <wchar.h>49932 **CHANGE HISTORY**49933 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
49934 (E).49935 **Issue 6**

49936 In the DESCRIPTION, a note on using this function in a threaded application is added.

49937 Extensions beyond the ISO C standard are now marked.

49938 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

49939 The *wcsrombs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

49940 **NAME**49941 `wcsspnp` — get length of a wide substring49942 **SYNOPSIS**49943 `#include <wchar.h>`49944 `size_t wcsspnp(const wchar_t *ws1, const wchar_t *ws2);`49945 **DESCRIPTION**

49946 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49947 conflict between the requirements described here and the ISO C standard is unintentional. This  
49948 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

49949 The `wcsspnp()` function shall compute the length of the maximum initial segment of the wide-  
49950 character string pointed to by `ws1` which consists entirely of wide-character codes from the  
49951 wide-character string pointed to by `ws2`.

49952 **RETURN VALUE**

49953 The `wcsspnp()` function shall return the length of the initial substring of `ws1`; no return value is  
49954 reserved to indicate an error.

49955 **ERRORS**

49956 No errors are defined.

49957 **EXAMPLES**

49958 None.

49959 **APPLICATION USAGE**

49960 None.

49961 **RATIONALE**

49962 None.

49963 **FUTURE DIRECTIONS**

49964 None.

49965 **SEE ALSO**49966 `wcscspnp()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`49967 **CHANGE HISTORY**

49968 First released in Issue 4. Derived from the MSE working draft.

49969 **Issue 5**

49970 The RETURN VALUE section is updated to indicate that `wcsspnp()` returns the length of `ws1`  
49971 rather than `ws1` itself.

49972 **NAME**49973 `wcsstr` — find a wide-character substring49974 **SYNOPSIS**49975 `#include <wchar.h>`49976 `wchar_t *wcsstr(const wchar_t *restrict ws1, const wchar_t *restrict ws2);`49977 **DESCRIPTION**49978 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
49979 conflict between the requirements described here and the ISO C standard is unintentional. This  
49980 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.49981 The `wcsstr()` function shall locate the first occurrence in the wide-character string pointed to by  
49982 `ws1` of the sequence of wide characters (excluding the terminating null wide character) in the  
49983 wide-character string pointed to by `ws2`.49984 **RETURN VALUE**49985 Upon successful completion, `wcsstr()` shall return a pointer to the located wide-character string,  
49986 or a null pointer if the wide-character string is not found.49987 If `ws2` points to a wide-character string with zero length, the function shall return `ws1`.49988 **ERRORS**

49989 No errors are defined.

49990 **EXAMPLES**

49991 None.

49992 **APPLICATION USAGE**

49993 None.

49994 **RATIONALE**

49995 None.

49996 **FUTURE DIRECTIONS**

49997 None.

49998 **SEE ALSO**49999 `wcschr()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`50000 **CHANGE HISTORY**50001 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50002 (E).50003 **Issue 6**50004 The `wcsstr()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

## 50005 NAME

50006 `wcstod`, `wcstof`, `wcstold` — convert a wide-character string to a double-precision number

## 50007 SYNOPSIS

50008 `#include <wchar.h>`

50009 `double wcstod(const wchar_t *restrict nptr, wchar_t **restrict endptr);`

50010 `float wcstof(const wchar_t *restrict nptr, wchar_t **restrict endptr);`

50011 `long double wcstold(const wchar_t *restrict nptr,`

50012 `wchar_t **restrict endptr);`

## 50013 DESCRIPTION

50014 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50015 conflict between the requirements described here and the ISO C standard is unintentional. This  
50016 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50017 These functions shall convert the initial portion of the wide-character string pointed to by *nptr* to  
50018 **double**, **float**, and **long double** representation, respectively. First, they decompose the input  
50019 wide-character string into three parts:

- 50020 1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by  
50021 `iswspace()`)
- 50022 2. A subject sequence interpreted as a floating-point constant or representing infinity or NaN
- 50023 3. A final wide-character string of one or more unrecognized wide-character codes, including  
50024 the terminating null wide-character code of the input wide-character string

50025 Then it attempts to convert the subject sequence to a floating-point number, and returns the  
50026 result.

50027 The expected form of the subject sequence is an optional plus or minus sign, then one of the  
50028 following:

- 50029 • A non-empty sequence of decimal digits optionally containing a radix character, then an  
50030 optional exponent part
- 50031 • A 0x or 0X, then a non-empty sequence of hexadecimal digits optionally containing a radix  
50032 character, then an optional binary exponent part
- 50033 • One of INF or INFINITY, or any other wide string equivalent except for case
- 50034 • One of NAN or NAN(*n-wchar-sequence<sub>opt</sub>*), or any other wide string ignoring case in the NAN  
50035 part, where:

50036 `n-wchar-sequence:`

50037 `digit`

50038 `nondigit`

50039 `n-wchar-sequence digit`

50040 `n-wchar-sequence nondigit`

50041 The subject sequence is defined as the longest initial subsequence of the input wide string,  
50042 starting with the first non-white-space wide character, that is of the expected form. The subject  
50043 sequence contains no wide characters if the input wide string is not of the expected form.

50044 If the subject sequence has the expected form for a floating-point number, the sequence of wide  
50045 characters starting with the first digit or the radix character (whichever occurs first) is  
50046 interpreted as a floating constant according to the rules of the C language, except that the radix  
50047 character is used in place of a period, and that if neither an exponent part nor a radix character  
50048 appears in a decimal floating-point number, or if a binary exponent part does not appear in a

50049 hexadecimal floating-point number, an exponent part of the appropriate type with value zero is  
 50050 assumed to follow the last digit in the string. If the subject sequence begins with a minus sign,  
 50051 the sequence is interpreted as negated. A wide-character sequence INF or INFINITY is  
 50052 interpreted as an infinity, if representable in the return type, else like a floating constant that is  
 50053 too large for the range of the return type. A wide-character sequence NAN or NAN(*n-wchar-*  
 50054 *sequence<sub>opt</sub>*) is interpreted as a quiet NaN, if supported in the return type, else like a subject  
 50055 sequence part that does not have the expected form; the meaning of the *n-wchar* sequences is  
 50056 implementation-defined. A pointer to the final wide string is stored in the object pointed to by  
 50057 *endptr*, provided that *endptr* is not a null pointer.

50058 If the subject sequence has the hexadecimal form and FLT\_RADIX is a power of 2, the value  
 50059 resulting from the conversion is correctly rounded.

50060 CX The radix character is defined in the program's locale (category *LC\_NUMERIC*). In the POSIX  
 50061 locale, or in a locale where the radix character is not defined, the radix character shall default to a  
 50062 period ( '.' ).

50063 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be  
 50064 accepted.

50065 If the subject sequence is empty or does not have the expected form, no conversion is performed;  
 50066 the value of *nptr* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null  
 50067 pointer.

50068 The *wcstod()* function shall not change the setting of *errno* if successful.

50069 Because 0 is returned on error and is also a valid return on success, an application wishing to  
 50070 check for error situations should set *errno* to 0, then call *wcstod()*, then check *errno*.

#### 50071 RETURN VALUE

50072 Upon successful completion, these functions shall return the converted value. If no conversion  
 50073 could be performed, 0 shall be returned and *errno* may be set to [EINVAL].

50074 If the correct value is outside the range of representable values, HUGE\_VAL, HUGE\_VALF, or  
 50075 HUGE\_VALL shall be returned (according to the sign of the value), and *errno* shall be set to  
 50076 [ERANGE].

50077 If the correct value would cause underflow, a value whose magnitude is no greater than the  
 50078 smallest normalized positive number in the return type shall be returned and *errno* set to  
 50079 [ERANGE].

#### 50080 ERRORS

50081 The *wcstod()* function shall fail if:

50082 [ERANGE] The value to be returned would cause overflow or underflow.

50083 The *wcstod()* function may fail if:

50084 CX [EINVAL] No conversion could be performed.

50085 **EXAMPLES**

50086 None.

50087 **APPLICATION USAGE**

50088 If the subject sequence has the hexadecimal form and FLT\_RADIX is not a power of 2, the result  
 50089 should be one of the two numbers in the appropriate internal format that are adjacent to the  
 50090 hexadecimal floating source value, with the extra stipulation that the error should have a correct  
 50091 sign for the current rounding direction.

50092 If the subject sequence has the decimal form and at most DECIMAL\_DIG (defined in <float.h>)  
 50093 significant digits, the result should be correctly rounded. If the subject sequence *D* has the  
 50094 decimal form and more than DECIMAL\_DIG significant digits, consider the two bounding,  
 50095 adjacent decimal strings *L* and *U*, both having DECIMAL\_DIG significant digits, such that the  
 50096 values of *L*, *D*, and *U* satisfy " $L \leq D \leq U$ ". The result should be one of the (equal or  
 50097 adjacent) values that would be obtained by correctly rounding *L* and *U* according to the current  
 50098 rounding direction, with the extra stipulation that the error with respect to *D* should have a  
 50099 correct sign for the current rounding direction.

50100 **RATIONALE**

50101 None.

50102 **FUTURE DIRECTIONS**

50103 None.

50104 **SEE ALSO**

50105 *iswspace()*, *localeconv()*, *scanf()*, *setlocale()*, *wcstol()*, the Base Definitions volume of  
 50106 IEEE Std. 1003.1-200x, <float.h>, <wchar.h>, the Base Definitions volume of  
 50107 IEEE Std. 1003.1-200x, Chapter 7, Locale

50108 **CHANGE HISTORY**

50109 First released in Issue 4. Derived from the MSE working draft.

50110 **Issue 5**50111 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.50112 **Issue 6**

50113 Extensions beyond the ISO C standard are now marked.

50114 The following new requirements on POSIX implementations derive from alignment with the  
 50115 Single UNIX Specification:

50116 • In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
 50117 added if no conversion could be performed.

50118 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

50119 • The *wcstod()* prototype is updated.50120 • The *wcstof()* and *wcstold()* functions are added.

50121 • The DESCRIPTION, RETURN VALUE, and APPLICATION USAGE sections are extensively  
 50122 updated.

50123 **NAME**

50124           wcstoimax, wcstoumax — convert wide-character string to integer type

50125 **SYNOPSIS**

50126           #include <stddef.h>

50127           #include <inttypes.h>

50128           intmax\_t wcstoimax(const wchar\_t \*restrict nptr,

50129                            wchar\_t \*\*restrict endpnr, int base);

50130           uintmax\_t wcstoumax(const wchar\_t \*restrict nptr,

50131                            wchar\_t \*\*restrict endpnr, int base);

50132 **DESCRIPTION**

50133 **CX**       The functionality described on this reference page is aligned with the ISO C standard. Any  
50134 conflict between the requirements described here and the ISO C standard is unintentional. This  
50135 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50136       These functions shall be equivalent to the *wcstol()*, *wcstoll()*, *wcstoul()*, and *wcstoull()* functions,  
50137 respectively, except that the initial portion of the wide string shall be converted to **intmax\_t** and  
50138 **uintmax\_t** representation, respectively.

50139 **RETURN VALUE**

50140       These functions shall return the converted value, if any.

50141       If no conversion could be performed, zero shall be returned. If the correct value is outside the  
50142 range of representable values, {INTMAX\_MAX}, {INTMAX\_MIN}, or {UINTMAX\_MAX} shall  
50143 be returned (according to the return type and sign of the value, if any), and *errno* shall be set to  
50144 [ERANGE].

50145 **ERRORS**

50146       These functions shall fail if:

50147       [EINVAL]       The value of *base* is not supported.

50148       [ERANGE]      The value to be returned is not representable.

50149       These functions may fail if:

50150       [EINVAL]      No conversion could be performed.

50151 **EXAMPLES**

50152       None.

50153 **APPLICATION USAGE**

50154       None.

50155 **RATIONALE**

50156       None.

50157 **FUTURE DIRECTIONS**

50158       None.

50159 **SEE ALSO**

50160       *wcstol()*, *wcstoul()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <inttypes.h>,  
50161 <stddef.h>

50162 **CHANGE HISTORY**

50163       First released in Issue 6. Derived from the ISO/IEC 9899: 1999 standard.

50164 **NAME**50165 `wcstok` — split wide-character string into tokens50166 **SYNOPSIS**50167 `#include <wchar.h>`50168 `wchar_t *wcstok(wchar_t *restrict ws1, const wchar_t *restrict ws2,`  
50169 `wchar_t **restrict ptr);`50170 **DESCRIPTION**50171 **cx** The functionality described on this reference page is aligned with the ISO C standard. Any  
50172 conflict between the requirements described here and the ISO C standard is unintentional. This  
50173 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.50174 A sequence of calls to `wcstok()` breaks the wide-character string pointed to by `ws1` into a  
50175 sequence of tokens, each of which is delimited by a wide-character code from the wide-character  
50176 string pointed to by `ws2`. The third argument points to a caller-provided `wchar_t` pointer into  
50177 which the `wcstok()` function stores information necessary for it to continue scanning the same  
50178 wide-character string.50179 The first call in the sequence has `ws1` as its first argument, and is followed by calls with a null  
50180 pointer as their first argument. The separator string pointed to by `ws2` may be different from call  
50181 to call.50182 The first call in the sequence searches the wide-character string pointed to by `ws1` for the first  
50183 wide-character code that is *not* contained in the current separator string pointed to by `ws2`. If no  
50184 such wide-character code is found, then there are no tokens in the wide-character string pointed  
50185 to by `ws1` and `wcstok()` returns a null pointer. If such a wide-character code is found, it is the start  
50186 of the first token.50187 The `wcstok()` function then searches from there for a wide-character code that *is* contained in the  
50188 current separator string. If no such wide-character code is found, the current token extends to  
50189 the end of the wide-character string pointed to by `ws1`, and subsequent searches for a token shall  
50190 return a null pointer. If such a wide-character code is found, it is overwritten by a null wide-  
50191 character, which terminates the current token. The `wcstok()` function saves a pointer to the  
50192 following wide-character code, from which the next search for a token shall start.50193 Each subsequent call, with a null pointer as the value of the first argument, starts searching from  
50194 the saved pointer and behaves as described above.50195 The implementation shall behave as if no function calls `wcstok()`.50196 **RETURN VALUE**50197 Upon successful completion, the `wcstok()` function shall return a pointer to the first wide-  
50198 character code of a token. Otherwise, if there is no token, `wcstok()` shall return a null pointer.50199 **ERRORS**

50200 No errors are defined.

50201 **EXAMPLES**

50202           None.

50203 **APPLICATION USAGE**

50204           None.

50205 **RATIONALE**

50206           None.

50207 **FUTURE DIRECTIONS**

50208           None.

50209 **SEE ALSO**50210           The Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>50211 **CHANGE HISTORY**

50212           First released in Issue 4.

50213 **Issue 5**50214           Aligned with ISO/IEC 9899:1990/Amendment 1:1995 (E). Specifically, a third argument is  
50215           added to the definition of this function in the SYNOPSIS.50216 **Issue 6**50217           The *wcstok()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

## 50218 NAME

50219 wcstol, wcstoll — convert a wide-character string to a long integer

## 50220 SYNOPSIS

50221 #include &lt;wchar.h&gt;

50222 long wcstol(const wchar\_t \*restrict *nptr*, wchar\_t \*\*restrict *endptr*,  
50223 int *base*);50224 long long wcstoll(const wchar\_t \*restrict *nptr*,  
50225 wchar\_t \*\*restrict *endptr*, int *base*);

## 50226 DESCRIPTION

50227 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50228 conflict between the requirements described here and the ISO C standard is unintentional. This  
50229 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.50230 These functions shall convert the initial portion of the wide-character string pointed to by *nptr* to  
50231 **long**, **long long**, **unsigned long**, and **unsigned long long** representation, respectively. First, they  
50232 decompose the input string into three parts:

- 50233 1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by  
50234 *iswspace()*)
- 50235 2. A subject sequence interpreted as an integer represented in some radix determined by the  
50236 value of *base*
- 50237 3. A final wide-character string of one or more unrecognized wide-character codes, including  
50238 the terminating null wide-character code of the input wide-character string

50239 Then it attempts to convert the subject sequence to an integer, and returns the result.

50240 If *base* is 0, the expected form of the subject sequence is that of a decimal constant, octal constant,  
50241 or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A decimal  
50242 constant begins with a non-zero digit, and consists of a sequence of decimal digits. An octal  
50243 constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to '7'  
50244 only. A hexadecimal constant consists of the prefix "0x" or "0X" followed by a sequence of the  
50245 decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.50246 If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence  
50247 of letters and digits representing an integer with the radix specified by *base*, optionally preceded  
50248 by a '+' or '-' sign, but not including an integer suffix. The letters from 'a' (or 'A') to 'z'  
50249 (or 'Z') inclusive are ascribed the values 10 to 35; only letters whose ascribed values are less  
50250 than that of *base* are permitted. If the value of *base* is 16, the wide-character code representations  
50251 of "0x" or "0X" may optionally precede the sequence of letters and digits, following the sign if  
50252 present.50253 The subject sequence is defined as the longest initial subsequence of the input wide-character  
50254 string, starting with the first non-white-space wide-character code that is of the expected form.  
50255 The subject sequence contains no wide-character codes if the input wide-character string is  
50256 empty or consists entirely of white-space wide-character code, or if the first non-white-space  
50257 wide-character code is other than a sign or a permissible letter or digit.50258 If the subject sequence has the expected form and *base* is 0, the sequence of wide-character codes  
50259 starting with the first digit is interpreted as an integer constant. If the subject sequence has the  
50260 expected form and the value of *base* is between 2 and 36, it is used as the base for conversion,  
50261 ascribing to each letter its value as given above. If the subject sequence begins with a minus sign,  
50262 the value resulting from the conversion is negated. A pointer to the final wide-character string is  
50263 stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

- 50264 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be  
50265 accepted.
- 50266 If the subject sequence is empty or does not have the expected form, no conversion is performed;  
50267 the value of *nptr* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null  
50268 pointer.
- 50269 The *wcstol()* function shall not change the setting of *errno* if successful.
- 50270 Because 0, {LONG\_MIN} or {LLONG\_MIN} and {LONG\_MAX} or {LLONG\_MAX} are returned  
50271 on error and are also valid returns on success, an application wishing to check for error  
50272 situations should set *errno* to 0, then call *wcstol()*, then check *errno*.
- 50273 **RETURN VALUE**
- 50274 Upon successful completion, these functions shall return the converted value, if any. If no  
50275 CX conversion could be performed, 0 shall be returned and *errno* may be set to indicate the error. If  
50276 the correct value is outside the range of representable values, {LONG\_MAX} or {LONG\_MIN}  
50277 shall be returned (according to the sign of the value), and *errno* set to [ERANGE].
- 50278 **ERRORS**
- 50279 These functions shall fail if:
- 50280 CX [EINVAL] The value of *base* is not supported.
- 50281 [ERANGE] The value to be returned is not representable.
- 50282 These functions may fail if:
- 50283 CX [EINVAL] No conversion could be performed.
- 50284 **EXAMPLES**
- 50285 None.
- 50286 **APPLICATION USAGE**
- 50287 None.
- 50288 **RATIONALE**
- 50289 None.
- 50290 **FUTURE DIRECTIONS**
- 50291 None.
- 50292 **SEE ALSO**
- 50293 *iswalph()*, *scanf()*, *wcstod()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>
- 50294 **CHANGE HISTORY**
- 50295 First released in Issue 4. Derived from the MSE working draft.
- 50296 **Issue 5**
- 50297 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.
- 50298 **Issue 6**
- 50299 Extensions beyond the ISO C standard are now marked.
- 50300 The following new requirements on POSIX implementations derive from alignment with the  
50301 Single UNIX Specification:
- 50302 • In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is
  - 50303 added if no conversion could be performed.
- 50304 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

50305

- The *wcstol()* prototype is updated.

50306

- The *wcstoll()* function is added.

50307 **NAME**

50308 wcstombs — convert a wide-character string to a character string

50309 **SYNOPSIS**

50310 #include &lt;stdlib.h&gt;

50311 size\_t wcstombs(char \*restrict s, const wchar\_t \*restrict pwcs,  
50312 size\_t n);50313 **DESCRIPTION**50314 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
50315 conflict between the requirements described here and the ISO C standard is unintentional. This  
50316 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.50317 The *wcstombs()* function shall convert the sequence of wide-character codes that are in the array  
50318 pointed to by *pwcs* into a sequence of characters that begins in the initial shift state and stores  
50319 these characters into the array pointed to by *s*, stopping if a character would exceed the limit of *n*  
50320 total bytes or if a null byte is stored. Each wide-character code is converted as if by a call to  
50321 *wctomb()*, except that the shift state of *wctomb()* is not affected.50322 The behavior of this function shall be affected by the *LC\_CTYPE* category of the current locale.50323 No more than *n* bytes shall be modified in the array pointed to by *s*. If copying takes place  
50324 **CX** between objects that overlap, the behavior is undefined. If *s* is a null pointer, *wcstombs()* shall  
50325 return the length required to convert the entire array regardless of the value of *n*, but no values  
50326 are stored.50327 The *wcstombs()* function need not be reentrant. A function that is not required to be reentrant is  
50328 not required to be thread-safe.50329 **RETURN VALUE**50330 If a wide-character code is encountered that does not correspond to a valid character (of one or  
50331 more bytes each), *wcstombs()* shall return (**size\_t**)−1. Otherwise, *wcstombs()* shall return the  
50332 number of bytes stored in the character array, not including any terminating null byte. The array  
50333 shall not be null-terminated if the value returned is *n*.50334 **ERRORS**50335 The *wcstombs()* function may fail if:50336 **CX** [EILSEQ] A wide-character code does not correspond to a valid character.50337 **EXAMPLES**

50338 None.

50339 **APPLICATION USAGE**

50340 None.

50341 **RATIONALE**

50342 None.

50343 **FUTURE DIRECTIONS**

50344 None.

50345 **SEE ALSO**50346 *mblen()*, *mbtowc()*, *mbstowcs()*, *wctomb()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
50347 <stdlib.h>

50348 **CHANGE HISTORY**

50349 First released in Issue 4. Derived from the ISO C standard. |

50350 **Issue 6**

50351 The following new requirements on POSIX implementations derive from alignment with the  
50352 Single UNIX Specification:

- 50353 • The DESCRIPTION states the effect of when *s* is a null pointer.
- 50354 • The [EILSEQ] error condition is added.

50355 The *wcstombs()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard. |

## 50356 NAME

50357 wcstoul, wcstoull — convert a wide-character string to an unsigned long

## 50358 SYNOPSIS

50359 #include <wchar.h>

50360 long wcstoul(const wchar\_t \*restrict *nptr*, wchar\_t \*\*restrict *endptr*,  
50361 int *base*);

50362 long long wcstoull(const wchar\_t \*restrict *nptr*,  
50363 wchar\_t \*\*restrict *endptr*, int *base*);

## 50364 DESCRIPTION

50365 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50366 conflict between the requirements described here and the ISO C standard is unintentional. This  
50367 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50368 These functions shall convert the initial portion of the wide-character string pointed to by *nptr* to  
50369 **long**, **long long**, **unsigned long**, and **unsigned long long** representation, respectively. First, they  
50370 decompose the input wide-character string into three parts:

- 50371 1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by  
50372 *iswspace()*)
- 50373 2. A subject sequence interpreted as an integer represented in some radix determined by the  
50374 value of *base*
- 50375 3. A final wide-character string of one or more unrecognized wide-character codes, including  
50376 the terminating null wide-character code of the input wide-character string

50377 Then it attempts to convert the subject sequence to an unsigned integer, and returns the result.

50378 If *base* is 0, the expected form of the subject sequence is that of a decimal constant, octal constant,  
50379 or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A decimal  
50380 constant begins with a non-zero digit, and consists of a sequence of decimal digits. An octal  
50381 constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to '7'  
50382 only. A hexadecimal constant consists of the prefix "0x" or "0X" followed by a sequence of the  
50383 decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.

50384 If the value of *base* is between 2 and 36, the expected form of the subject sequence is a sequence  
50385 of letters and digits representing an integer with the radix specified by *base*, optionally preceded  
50386 by a '+' or '-' sign, but not including an integer suffix. The letters from 'a' (or 'A') to 'z'  
50387 (or 'Z') inclusive are ascribed the values 10 to 35; only letters whose ascribed values are less  
50388 than that of *base* are permitted. If the value of *base* is 16, the wide-character codes "0x" or "0X"  
50389 may optionally precede the sequence of letters and digits, following the sign if present.

50390 The subject sequence is defined as the longest initial subsequence of the input wide-character  
50391 string, starting with the first wide-character code that is not white space and is of the expected  
50392 form. The subject sequence contains no wide-character codes if the input wide-character string is  
50393 empty or consists entirely of white-space wide-character codes, or if the first wide-character  
50394 code that is not white space is other than a sign or a permissible letter or digit.

50395 If the subject sequence has the expected form and *base* is 0, the sequence of wide-character codes  
50396 starting with the first digit is interpreted as an integer constant. If the subject sequence has the  
50397 expected form and the value of *base* is between 2 and 36, it is used as the base for conversion,  
50398 ascribing to each letter its value as given above. If the subject sequence begins with a minus sign,  
50399 the value resulting from the conversion is negated. A pointer to the final wide-character string is  
50400 stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

50401 CX In other than the C or POSIX locales, other implementation-defined subject sequences may be  
50402 accepted.

50403 If the subject sequence is empty or does not have the expected form, no conversion is performed;  
50404 the value of *nptr* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null  
50405 pointer.

50406 The *wcstoul()* function shall not change the setting of *errno* if successful.

50407 Because 0, {ULONG\_MAX}, and {ULLONG\_MAX} are returned on error and 0 is also a valid  
50408 return on success, an application wishing to check for error situations should set *errno* to 0, then  
50409 call *wcstoul()*, then check *errno*.

#### 50410 RETURN VALUE

50411 Upon successful completion, these functions shall return the converted value, if any. If no  
50412 CX conversion could be performed, 0 shall be returned and *errno* may be set to indicate the error. If  
50413 the correct value is outside the range of representable values, {ULONG\_MAX} shall be returned  
50414 and *errno* set to [ERANGE].

#### 50415 ERRORS

50416 These functions shall fail if:

50417 CX [EINVAL] The value of *base* is not supported.

50418 [ERANGE] The value to be returned is not representable.

50419 These functions may fail if:

50420 CX [EINVAL] No conversion could be performed.

#### 50421 EXAMPLES

50422 None.

#### 50423 APPLICATION USAGE

50424 None.

#### 50425 RATIONALE

50426 None.

#### 50427 FUTURE DIRECTIONS

50428 None.

#### 50429 SEE ALSO

50430 *iswalph()*, *scanf()*, *wcstod()*, *wcstol()*, the Base Definitions volume of IEEE Std. 1003.1-200x,  
50431 <wchar.h>

#### 50432 CHANGE HISTORY

50433 First released in Issue 4. Derived from the MSE working draft.

#### 50434 Issue 5

50435 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.

#### 50436 Issue 6

50437 Extensions beyond the ISO C standard are now marked.

50438 The following new requirements on POSIX implementations derive from alignment with the  
50439 Single UNIX Specification:

- 50440 • The [EINVAL] error condition is added for when the value of *base* is not supported.

50441 In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
50442 added if no conversion could be performed.

50443 The following changes are made for alignment with the ISO/IEC 9899:1999 standard:

50444 • The *wcstoul()* prototype is updated.

50445 • The *wcstoull()* function is added.

50446 **NAME**50447           wcswcs — find a wide substring (**LEGACY**)50448 **SYNOPSIS**

50449 xSI       #include &lt;wchar.h&gt;

50450           wchar\_t \*wcswcs(const wchar\_t \*ws1, const wchar\_t \*ws2);

50451

50452 **DESCRIPTION**

50453           The `wcswcs()` function shall locate the first occurrence in the wide-character string pointed to by  
50454           `ws1` of the sequence of wide-character codes (excluding the terminating null wide-character  
50455           code) in the wide-character string pointed to by `ws2`.

50456 **RETURN VALUE**

50457           Upon successful completion, `wcswcs()` shall return a pointer to the located wide-character string  
50458           or a null pointer if the wide-character string is not found.

50459           If `ws2` points to a wide-character string with zero length, the function shall return `ws1`.

50460 **ERRORS**

50461           No errors are defined.

50462 **EXAMPLES**

50463           None.

50464 **APPLICATION USAGE**

50465           This function was not included in the final ISO/IEC 9899:1990/Amendment 1:1995 (E).

50466           Application developers are strongly encouraged to use the `wcsstr()` function instead.50467 **RATIONALE**

50468           None.

50469 **FUTURE DIRECTIONS**

50470           This function may be withdrawn in a future version.

50471 **SEE ALSO**50472           `wcschr()`, `wcsstr()`, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>50473 **CHANGE HISTORY**

50474           First released in Issue 4. Derived from the MSE working draft.

50475 **Issue 5**

50476           Marked EX.

50477 **Issue 6**

50478           This function is marked LEGACY.

50479 **NAME**

50480 `wcswidth` — number of column positions of a wide-character string

50481 **SYNOPSIS**

```
50482 xsi #include <wchar.h>
```

```
50483 int wcswidth(const wchar_t *pwcs, size_t n);
```

50484

50485 **DESCRIPTION**

50486 The `wcswidth()` function shall determine the number of column positions required for  $n$  wide-  
50487 character codes (or fewer than  $n$  wide-character codes if a null wide-character code is  
50488 encountered before  $n$  wide-character codes are exhausted) in the string pointed to by `pwcs`.

50489 **RETURN VALUE**

50490 The `wcswidth()` function either shall return 0 (if `pwcs` points to a null wide-character code), or  
50491 return the number of column positions to be occupied by the wide-character string pointed to by  
50492 `pwcs`, or return -1 (if any of the first  $n$  wide-character codes in the wide-character string pointed  
50493 to by `pwcs` is not a printing wide-character code).

50494 **ERRORS**

50495 No errors are defined.

50496 **EXAMPLES**

50497 None.

50498 **APPLICATION USAGE**

50499 This function was removed from the final ISO/IEC 9899:1990/Amendment 1:1995 (E), and the  
50500 return value for a non-printable wide character is not specified.

50501 **RATIONALE**

50502 None.

50503 **FUTURE DIRECTIONS**

50504 None.

50505 **SEE ALSO**

50506 `wcwidth()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`, the Base Definitions  
50507 volume of IEEE Std. 1003.1-200x, Section 3.106, Column Position

50508 **CHANGE HISTORY**

50509 First released in Issue 4. Derived from the MSE working draft.

50510 **Issue 6**

50511 The Open Group corrigenda item U021/11 has been applied. The function is marked as an  
50512 extension.

50513 **NAME**50514        `wcsxfrm` — wide-character string transformation50515 **SYNOPSIS**50516        `#include <wchar.h>`50517        `size_t wcsxfrm(wchar_t *restrict ws1, const wchar_t *restrict ws2,`  
50518            `size_t n);`50519 **DESCRIPTION**50520 **CX**        The functionality described on this reference page is aligned with the ISO C standard. Any  
50521 conflict between the requirements described here and the ISO C standard is unintentional. This  
50522 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.50523        The `wcsxfrm()` function shall transform the wide-character string pointed to by `ws2` and place the  
50524 resulting wide-character string into the array pointed to by `ws1`. The transformation is such that  
50525 if `wscmp()` is applied to two transformed wide strings, it returns a value greater than, equal to,  
50526 or less than 0, corresponding to the result of `wscoll()` applied to the same two original wide-  
50527 character strings. No more than `n` wide-character codes are placed into the resulting array  
50528 pointed to by `ws1`, including the terminating null wide-character code. If `n` is 0, `ws1` is permitted  
50529 to be a null pointer. If copying takes place between objects that overlap, the behavior is  
50530 undefined.50531 **CX**        The `wcsxfrm()` function shall not change the setting of `errno` if successful.50532        Because no return value is reserved to indicate an error, an application wishing to check for error  
50533 situations should set `errno` to 0, then call `wcsxfrm()`, then check `errno`.50534 **RETURN VALUE**50535        The `wcsxfrm()` function shall return the length of the transformed wide-character string (not  
50536 including the terminating null wide-character code). If the value returned is `n` or more, the  
50537 contents of the array pointed to by `ws1` are indeterminate.50538        On error, the `wcsxfrm()` function may set `errno`, but no return value is reserved to indicate an  
50539 error.50540 **ERRORS**50541        The `wcsxfrm()` function may fail if:50542 **CX**        **[EINVAL]**        The wide-character string pointed to by `ws2` contains wide-character codes  
50543 outside the domain of the collating sequence.50544 **EXAMPLES**

50545        None.

50546 **APPLICATION USAGE**50547        The transformation function is such that two transformed wide-character strings can be ordered  
50548 by `wscmp()` as appropriate to collating sequence information in the program's locale (category  
50549 `LC_COLLATE`).50550        The fact that when `n` is 0 `ws1` is permitted to be a null pointer is useful to determine the size of  
50551 the `ws1` array prior to making the transformation.50552 **RATIONALE**

50553        None.

50554 **FUTURE DIRECTIONS**

50555 None.

50556 **SEE ALSO**50557 `wscmp()`, `wscoll()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`50558 **CHANGE HISTORY**

50559 First released in Issue 4. Derived from the MSE working draft.

50560 **Issue 5**

50561 Moved from ENHANCED I18N to BASE and the [ENOSYS] error is removed.

50562 The DESCRIPTION is updated to indicate that *errno* is not changed if the function is successful.50563 **Issue 6**50564 In previous versions, this function was required to return `-1` on error.

50565 Extensions beyond the ISO C standard are now marked.

50566 The following new requirements on POSIX implementations derive from alignment with the  
50567 Single UNIX Specification:

- 50568
- In the RETURN VALUE and ERRORS sections, the [EINVAL] optional error condition is  
50569 added if no conversion could be performed.

50570 The `wcsxfrm()` prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

50571 **NAME**

50572 wctob — wide-character to single-byte conversion

50573 **SYNOPSIS**

50574 #include &lt;stdio.h&gt;

50575 #include &lt;wchar.h&gt;

50576 int wctob(wint\_t c);

50577 **DESCRIPTION**

50578 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50579 conflict between the requirements described here and the ISO C standard is unintentional. This  
50580 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50581 The *wctob()* function shall determine whether *c* corresponds to a member of the extended  
50582 character set whose character representation is a single byte when in the initial shift state.

50583 The behavior of this function shall be affected by the *LC\_CTYPE* category of the current locale.

50584 **RETURN VALUE**

50585 The *wctob()* function shall return EOF if *c* does not correspond to a character with length one in  
50586 the initial shift state. Otherwise, it shall return the single-byte representation of that character as  
50587 an **unsigned char** converted to **int**.

50588 **ERRORS**

50589 No errors are defined.

50590 **EXAMPLES**

50591 None.

50592 **APPLICATION USAGE**

50593 None.

50594 **RATIONALE**

50595 None.

50596 **FUTURE DIRECTIONS**

50597 None.

50598 **SEE ALSO**50599 *btowc()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <**wchar.h**>50600 **CHANGE HISTORY**

50601 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50602 (E).

50603 **NAME**

50604 wctomb — convert a wide-character code to a character

50605 **SYNOPSIS**

50606 #include &lt;stdlib.h&gt;

50607 int wctomb(char \*s, wchar\_t wchar);

50608 **DESCRIPTION**

50609 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50610 conflict between the requirements described here and the ISO C standard is unintentional. This  
50611 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50612 The *wctomb()* function shall determine the number of bytes needed to represent the character  
50613 corresponding to the wide-character code whose value is *wchar* (including any change in the  
50614 shift state). It stores the character representation (possibly multiple bytes and any special bytes  
50615 to change shift state) in the array object pointed to by *s* (if *s* is not a null pointer). At most  
50616 {MB\_CUR\_MAX} bytes are stored. If *wchar* is 0, a null byte is stored, preceded by any shift  
50617 sequence needed to restore the initial shift state, and *wctomb()* is left in the initial shift state.

50618 cx The behavior of this function is affected by the *LC\_CTYPE* category of the current locale. For a  
50619 state-dependent encoding, this function is placed into its initial state by a call for which its  
50620 character pointer argument, *s*, is a null pointer. Subsequent calls with *s* as other than a null  
50621 pointer cause the internal state of the function to be altered as necessary. A call with *s* as a null  
50622 pointer causes this function to return a non-zero value if encodings have state dependency, and  
50623 0 otherwise. Changing the *LC\_CTYPE* category causes the shift state of this function to be  
50624 indeterminate.

50625 The *wctomb()* function need not be reentrant. A function that is not required to be reentrant is  
50626 not required to be thread-safe.

50627 The implementation shall behave as if no function defined in this volume of  
50628 IEEE Std. 1003.1-200x calls *wctomb()*.

50629 **RETURN VALUE**

50630 If *s* is a null pointer, *wctomb()* shall return a non-zero or 0 value, if character encodings,  
50631 respectively, do or do not have state-dependent encodings. If *s* is not a null pointer, *wctomb()*  
50632 shall return -1 if the value of *wchar* does not correspond to a valid character, or return the  
50633 number of bytes that constitute the character corresponding to the value of *wchar*.

50634 In no case shall the value returned be greater than the value of the {MB\_CUR\_MAX} macro.

50635 **ERRORS**

50636 No errors are defined.

50637 **EXAMPLES**

50638 None.

50639 **APPLICATION USAGE**

50640 None.

50641 **RATIONALE**

50642 None.

50643 **FUTURE DIRECTIONS**

50644 None.

50645 **SEE ALSO**

50646            *mblen()*, *mbtowc()*, *mbstowcs()*, *wcstombs()*, the Base Definitions volume of IEEE Std. 1003.1-200x, |  
50647            <stdlib.h>

50648 **CHANGE HISTORY**

50649            First released in Issue 4. Derived from the ANSI C standard. |

50650 **Issue 6**

50651            Extensions beyond the ISO C standard are now marked.

50652            In the DESCRIPTION, a note about reentrancy and thread-safety is added.

50653 **NAME**

50654 wctrans — define character mapping

50655 **SYNOPSIS**

50656 #include &lt;wctype.h&gt;

50657 wctrans\_t wctrans(const char \*charclass);

50658 **DESCRIPTION**

50659 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50660 conflict between the requirements described here and the ISO C standard is unintentional. This  
50661 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50662 The *wctrans()* function is defined for valid character mapping names identified in the current  
50663 locale. The *charclass* is a string identifying a generic character mapping name for which codeset-  
50664 specific information is required. The following character mapping names are defined in all  
50665 locales: **tolower** and **toupper**.

50666 The function shall return a value of type **wctrans\_t**, which can be used as the second argument  
50667 to subsequent calls of *towctrans()*. The *wctrans()* function determines values of **wctrans\_t**  
50668 according to the rules of the coded character set defined by character mapping information in  
50669 the program's locale (category *LC\_CTYPE*). The values returned by *wctrans()* are valid until a  
50670 call to *setlocale()* that modifies the category *LC\_CTYPE*.

50671 **RETURN VALUE**

50672 cx The *wctrans()* function shall return 0 and may set *errno* to indicate the error if the given character  
50673 mapping name is not valid for the current locale (category *LC\_CTYPE*); otherwise, it shall return  
50674 a non-zero object of type **wctrans\_t** that can be used in calls to *towctrans()*.

50675 **ERRORS**50676 The *wctrans()* function may fail if:

50677 cx [EINVAL] The character mapping name pointed to by *charclass* is not valid in the current  
50678 locale.

50679 **EXAMPLES**

50680 None.

50681 **APPLICATION USAGE**

50682 None.

50683 **RATIONALE**

50684 None.

50685 **FUTURE DIRECTIONS**

50686 None.

50687 **SEE ALSO**50688 *towctrans()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <wctype.h>50689 **CHANGE HISTORY**

50690 First released in Issue 5. Derived from ISO/IEC 9899:1990/Amendment 1:1995 (E).

50691 **NAME**

50692 wctype — define character class

50693 **SYNOPSIS**

50694 #include &lt;wctype.h&gt;

50695 wctype\_t wctype(const char \*property);

50696 **DESCRIPTION**

50697 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
 50698 conflict between the requirements described here and the ISO C standard is unintentional. This  
 50699 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50700 The *wctype()* function is defined for valid character class names as defined in the current locale.  
 50701 The *property* is a string identifying a generic character class for which codeset-specific type  
 50702 information is required. The following character class names are defined in all locales:

|       |              |              |               |
|-------|--------------|--------------|---------------|
| 50703 | <b>alnum</b> | <b>digit</b> | <b>punct</b>  |
| 50704 | <b>alpha</b> | <b>graph</b> | <b>space</b>  |
| 50705 | <b>blank</b> | <b>lower</b> | <b>upper</b>  |
| 50706 | <b>cntrl</b> | <b>print</b> | <b>xdigit</b> |

50707 Additional character class names defined in the locale definition file (category *LC\_CTYPE*) can  
 50708 also be specified.

50709 The function shall return a value of type **wctype\_t**, which can be used as the second argument to  
 50710 subsequent calls of *iswctype()*. The *wctype()* function determines values of **wctype\_t** according  
 50711 to the rules of the coded character set defined by character type information in the program's  
 50712 locale (category *LC\_CTYPE*). The values returned by *wctype()* are valid until a call to *setlocale()*  
 50713 that modifies the category *LC\_CTYPE*.

50714 **RETURN VALUE**

50715 The *wctype()* function shall return 0 if the given character class name is not valid for the current  
 50716 locale (category *LC\_CTYPE*); otherwise, it shall return an object of type **wctype\_t** that can be  
 50717 used in calls to *iswctype()*.

50718 **ERRORS**

50719 No errors are defined.

50720 **EXAMPLES**

50721 None.

50722 **APPLICATION USAGE**

50723 None.

50724 **RATIONALE**

50725 None.

50726 **FUTURE DIRECTIONS**

50727 None.

50728 **SEE ALSO**50729 *iswctype()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <wctype.h>, <wchar.h>50730 **CHANGE HISTORY**

50731 First released in Issue 4.

50732 **Issue 5**

50733 The following change has been made in this issue for alignment with  
50734 ISO/IEC 9899:1990/Amendment 1:1995 (E):

- 50735 • The SYNOPSIS has been changed to indicate that this function and associated data types are  
50736 now made visible by inclusion of the header `<wctype.h>` rather than `<wchar.h>`.

50737 **NAME**

50738 `wcwidth` — number of column positions of a wide-character code

50739 **SYNOPSIS**

```
50740 xSI #include <wchar.h>
```

```
50741 int wcwidth(wchar_t wc);
```

50742

50743 **DESCRIPTION**

50744 The `wcwidth()` function shall determine the number of column positions required for the wide  
50745 character `wc`. The application shall ensure that the value of `wc` is a character representable as a  
50746 **wchar\_t**, and a wide-character code corresponding to a valid character in the current locale.

50747 **RETURN VALUE**

50748 The `wcwidth()` function shall either return 0 (if `wc` is a null wide-character code), or return the  
50749 number of column positions to be occupied by the wide-character code `wc`, or return -1 (if `wc`  
50750 does not correspond to a printing wide-character code).

50751 **ERRORS**

50752 No errors are defined.

50753 **EXAMPLES**

50754 None.

50755 **APPLICATION USAGE**

50756 This function was removed from the final ISO/IEC 9899:1990/Amendment 1:1995 (E), and the  
50757 return value for a non-printable wide character is not specified.

50758 **RATIONALE**

50759 None.

50760 **FUTURE DIRECTIONS**

50761 None.

50762 **SEE ALSO**

50763 `wcswidth()`, the Base Definitions volume of IEEE Std. 1003.1-200x, `<wchar.h>`

50764 **CHANGE HISTORY**

50765 First released as a World-wide Portability Interface in Issue 4. Derived from MSE working draft.

50766 **Issue 6**

50767 The Open Group corrigenda item U021/12 has been applied. This function is marked as an  
50768 extension.

50769 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50770 **NAME**

50771 wmemchr — find a wide character in memory

50772 **SYNOPSIS**

50773 #include &lt;wchar.h&gt;

50774 wchar\_t \*wmemchr(const wchar\_t \*ws, wchar\_t wc, size\_t n);

50775 **DESCRIPTION**

50776 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50777 conflict between the requirements described here and the ISO C standard is unintentional. This  
50778 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50779 The *wmemchr()* function shall locate the first occurrence of *wc* in the initial *n* wide characters of  
50780 the object pointed to by *ws*. This function is not affected by locale and all **wchar\_t** values are  
50781 treated identically. The null wide character and **wchar\_t** values not corresponding to valid  
50782 characters are not treated specially.

50783 If *n* is zero, the application shall ensure that *ws* is a valid pointer and the function behaves as if  
50784 no valid occurrence of *wc* is found.

50785 **RETURN VALUE**

50786 The *wmemchr()* function shall return a pointer to the located wide character, or a null pointer if  
50787 the wide character does not occur in the object.

50788 **ERRORS**

50789 No errors are defined.

50790 **EXAMPLES**

50791 None.

50792 **APPLICATION USAGE**

50793 None.

50794 **RATIONALE**

50795 None.

50796 **FUTURE DIRECTIONS**

50797 None.

50798 **SEE ALSO**

50799 *wmemcmp()*, *wmemcpy()*, *wmemmove()*, *wmemset()*, the Base Definitions volume of  
50800 IEEE Std. 1003.1-200x, <wchar.h>

50801 **CHANGE HISTORY**

50802 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50803 (E).

50804 **Issue 6**

50805 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50806 **NAME**

50807 wmemcmp — compare wide characters in memory

50808 **SYNOPSIS**

50809 #include &lt;wchar.h&gt;

50810 int wmemcmp(const wchar\_t \*ws1, const wchar\_t \*ws2, size\_t n);

50811 **DESCRIPTION**

50812 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50813 conflict between the requirements described here and the ISO C standard is unintentional. This  
50814 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50815 The *wmemcmp()* function shall compare the first *n* wide characters of the object pointed to by  
50816 *ws1* to the first *n* wide characters of the object pointed to by *ws2*. This function is not affected by  
50817 locale and all **wchar\_t** values are treated identically. The null wide character and **wchar\_t** values  
50818 not corresponding to valid characters are not treated specially.

50819 If *n* is zero, the application shall ensure that *ws1* and *ws2* are valid pointers, and the function  
50820 behaves as if the two objects compare equal.

50821 **RETURN VALUE**

50822 The *wmemcmp()* function shall return an integer greater than, equal to, or less than zero,  
50823 respectively, as the object pointed to by *ws1* is greater than, equal to, or less than the object  
50824 pointed to by *ws2*.

50825 **ERRORS**

50826 No errors are defined.

50827 **EXAMPLES**

50828 None.

50829 **APPLICATION USAGE**

50830 None.

50831 **RATIONALE**

50832 None.

50833 **FUTURE DIRECTIONS**

50834 None.

50835 **SEE ALSO**

50836 *wmemchr()*, *wmemcpy()*, *wmemmove()*, *wmemset()*, the Base Definitions volume of  
50837 IEEE Std. 1003.1-200x, <wchar.h>

50838 **CHANGE HISTORY**

50839 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50840 (E).

50841 **Issue 6**

50842 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50843 **NAME**

50844 wmemcpy — copy wide characters in memory

50845 **SYNOPSIS**

50846 #include &lt;wchar.h&gt;

50847 wchar\_t \*wmemcpy(wchar\_t \*restrict *ws1*, const wchar\_t \*restrict *ws2*,  
50848 size\_t *n*);50849 **DESCRIPTION**50850 **CX** The functionality described on this reference page is aligned with the ISO C standard. Any  
50851 conflict between the requirements described here and the ISO C standard is unintentional. This  
50852 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.50853 The *wmemcpy()* function shall copy *n* wide characters from the object pointed to by *ws2* to the  
50854 object pointed to by *ws1*. This function is not affected by locale and all **wchar\_t** values are  
50855 treated identically. The null wide character and **wchar\_t** values not corresponding to valid  
50856 characters are not treated specially.50857 If *n* is zero, the application shall ensure that *ws1* and *ws2* are valid pointers, and the function  
50858 copies zero wide characters.50859 **RETURN VALUE**50860 The *wmemcpy()* function shall return the value of *ws1*.50861 **ERRORS**

50862 No errors are defined.

50863 **EXAMPLES**

50864 None.

50865 **APPLICATION USAGE**

50866 None.

50867 **RATIONALE**

50868 None.

50869 **FUTURE DIRECTIONS**

50870 None.

50871 **SEE ALSO**50872 *wmemchr()*, *wmemcmp()*, *wmemmove()*, *wmemset()*, the Base Definitions volume of  
50873 IEEE Std. 1003.1-200x, <wchar.h>50874 **CHANGE HISTORY**50875 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50876 (E).50877 **Issue 6**

50878 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50879 The *wmemcpy()* prototype is updated for alignment with the ISO/IEC 9899:1999 standard.

50880 **NAME**

50881 wmemmove — copy wide characters in memory with overlapping areas

50882 **SYNOPSIS**

50883 #include &lt;wchar.h&gt;

50884 wchar\_t \*wmemmove(wchar\_t \*ws1, const wchar\_t \*ws2, size\_t n);

50885 **DESCRIPTION**

50886 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50887 conflict between the requirements described here and the ISO C standard is unintentional. This  
50888 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50889 The *wmemmove()* function shall copy *n* wide characters from the object pointed to by *ws2* to the  
50890 object pointed to by *ws1*. Copying takes place as if the *n* wide characters from the object pointed  
50891 to by *ws2* are first copied into a temporary array of *n* wide characters that does not overlap the  
50892 objects pointed to by *ws1* or *ws2*, and then the *n* wide characters from the temporary array are  
50893 copied into the object pointed to by *ws1*.

50894 This function is not affected by locale and all **wchar\_t** values are treated identically. The null  
50895 wide character and **wchar\_t** values not corresponding to valid characters are not treated  
50896 specially.

50897 If *n* is zero, the application shall ensure that *ws1* and *ws2* are valid pointers, and the function  
50898 copies zero wide characters.

50899 **RETURN VALUE**50900 The *wmemmove()* function shall return the value of *ws1*.50901 **ERRORS**

50902 No errors are defined

50903 **EXAMPLES**

50904 None.

50905 **APPLICATION USAGE**

50906 None.

50907 **RATIONALE**

50908 None.

50909 **FUTURE DIRECTIONS**

50910 None.

50911 **SEE ALSO**

50912 *wmemchr()*, *wmemcmp()*, *wmemcpy()*, *wmemset()*, the Base Definitions volume of  
50913 IEEE Std. 1003.1-200x, <wchar.h>

50914 **CHANGE HISTORY**

50915 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50916 (E).

50917 **Issue 6**

50918 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50919 **NAME**

50920 wmemset — set wide characters in memory

50921 **SYNOPSIS**

50922 #include <wchar.h>

50923 wchar\_t \*wmemset(wchar\_t \*ws, wchar\_t wc, size\_t n);

50924 **DESCRIPTION**

50925 cx The functionality described on this reference page is aligned with the ISO C standard. Any  
50926 conflict between the requirements described here and the ISO C standard is unintentional. This  
50927 volume of IEEE Std. 1003.1-200x defers to the ISO C standard.

50928 The *wmemset()* function shall copy the value of *wc* into each of the first *n* wide characters of the  
50929 object pointed to by *ws*. This function is not affected by locale and all **wchar\_t** values are treated  
50930 identically. The null wide character and **wchar\_t** values not corresponding to valid characters  
50931 are not treated specially.

50932 If *n* is zero, the application shall ensure that *ws* is a valid pointer, and the function copies zero  
50933 wide characters.

50934 **RETURN VALUE**

50935 The *wmemset()* functions shall return the value of *ws*.

50936 **ERRORS**

50937 No errors are defined.

50938 **EXAMPLES**

50939 None.

50940 **APPLICATION USAGE**

50941 None.

50942 **RATIONALE**

50943 None.

50944 **FUTURE DIRECTIONS**

50945 None.

50946 **SEE ALSO**

50947 *wmemchr()*, *wmemcmp()*, *wmemcpy()*, *wmemmove()*, the Base Definitions volume of  
50948 IEEE Std. 1003.1-200x, <wchar.h>

50949 **CHANGE HISTORY**

50950 First released in Issue 5. Included for alignment with ISO/IEC 9899:1990/Amendment 1:1995  
50951 (E).

50952 **Issue 6**

50953 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

50954 **NAME**

50955 wordexp, wordfree — perform word expansions

50956 **SYNOPSIS**

50957 #include &lt;wordexp.h&gt;

50958 int wordexp(const char \*restrict words, wordexp\_t \*restrict pwordexp,  
50959 int flags);

50960 void wordfree(wordexp\_t \*pwordexp);

50961 **DESCRIPTION**50962 The *wordexp()* function shall perform word expansions and place the list of expanded words  
50963 into *pwordexp*.50964 If the implementation supports the utilities defined in the Shell and Utilities volume of  
50965 IEEE Std. 1003.1-200x, the *wordexp()* function performs word expansions as described in the  
50966 Shell and Utilities volume of IEEE Std. 1003.1-200x, Section 2.6, Word Expansions, subject to  
50967 quoting as in the Shell and Utilities volume of IEEE Std. 1003.1-200x, Section 2.2, Quoting, and  
50968 places the list of expanded words into the structure pointed to by *pwordexp*.50969 The *words* argument is a pointer to a string containing one or more words to be expanded. The  
50970 expansions shall be the same as would be performed by the command line interpreter if *words*  
50971 were the part of a command line representing the arguments to a utility. Therefore, the  
50972 application shall ensure that *words* does not contain an unquoted <newline> or any of the  
50973 unquoted shell special characters '|', '&', ';', '<', '>' except in the context of command  
50974 substitution as specified in the Shell and Utilities volume of IEEE Std. 1003.1-200x, Section 2.6.3,  
50975 Command Substitution. It also shall not contain unquoted parentheses or braces, except in the  
50976 context of command or variable substitution. If the argument *words* contains an unquoted  
50977 comment character (number sign) that is the beginning of a token, *wordexp()* shall either treat the  
50978 comment character as a regular character, or interpret it as a comment indicator and ignore the  
50979 remainder of *words*.50980 If the implementation does not support the utilities defined in the Shell and Utilities volume of  
50981 IEEE Std. 1003.1-200x, the word expansion is unspecified, but should be the same as that used by  
50982 the command language interpreter used by the *system()* and *popen()* functions.50983 The structure type **wordexp\_t** is defined in the header <**wordexp.h**> and includes at least the  
50984 following members:

50985

50986

50987

50988

50989

| Member Type | Member Name | Description                                                         |
|-------------|-------------|---------------------------------------------------------------------|
| size_t      | we_wordc    | Count of words matched by <i>words</i> .                            |
| char **     | we_wordv    | Pointer to list of expanded words.                                  |
| size_t      | we_offs     | Slots to reserve at the beginning of <i>pwordexp-&gt;we_wordv</i> . |

50990 The *wordexp()* function stores the number of generated words into *pwordexp->we\_wordc* and a  
50991 pointer to a list of pointers to words in *pwordexp->we\_wordv*. If the implementation supports the  
50992 utilities defined in the Shell and Utilities volume of IEEE Std. 1003.1-200x, each individual field  
50993 created during field splitting (see the Shell and Utilities volume of IEEE Std. 1003.1-200x, Section  
50994 2.6.5, Field Splitting) or path name expansion (see the Shell and Utilities volume of  
50995 IEEE Std. 1003.1-200x, Section 2.6.6, Path Name Expansion) is a separate word in the *pwordexp-*  
50996 *>we\_wordv* list. The words are in order as described in the Shell and Utilities volume of  
50997 IEEE Std. 1003.1-200x, Section 2.6, Word Expansions. The first pointer after the last word pointer  
50998 shall be a null pointer. The expansion of special parameters described in the Shell and Utilities  
50999 volume of IEEE Std. 1003.1-200x, Section 2.5.2, Special Parameters is unspecified.

51000 It is the caller's responsibility to allocate the storage pointed to by *pwordexp*. The *wordexp()*  
 51001 function allocates other space as needed, including memory pointed to by *pwordexp->we\_wordv*.  
 51002 The *wordfree()* function frees any memory associated with *pwordexp* from a previous call to  
 51003 *wordexp()*.

51004 The *flags* argument is used to control the behavior of *wordexp()*. The value of *flags* is the  
 51005 bitwise-inclusive OR of zero or more of the following constants, which are defined in  
 51006 **<wordexp.h>**:

51007 **WRDE\_APPEND** Append words generated to the ones from a previous call to *wordexp()*.

51008 **WRDE\_DOOFFS** Make use of *pwordexp->we\_offs*. If this flag is set, *pwordexp->we\_offs* is used  
 51009 to specify how many null pointers to add to the beginning of *pwordexp->we\_wordv*. In other words,  
 51010 *pwordexp->we\_wordv* shall point to *pwordexp->we\_offs* null pointers, followed by *pwordexp->we\_wordc*  
 51011 word pointers, followed by a null pointer.  
 51012

51013 **WRDE\_NOCMD** If the implementation supports the utilities defined in the Shell and  
 51014 Utilities volume of IEEE Std. 1003.1-200x, fail if command substitution, as  
 51015 specified in the Shell and Utilities volume of IEEE Std. 1003.1-200x,  
 51016 Section 2.6.3, Command Substitution, is requested.

51017 **WRDE\_REUSE** The *pwordexp* argument was passed to a previous successful call to  
 51018 *wordexp()*, and has not been passed to *wordfree()*. The result shall be the  
 51019 same as if the application had called *wordfree()* and then called *wordexp()*  
 51020 without **WRDE\_REUSE**.

51021 **WRDE\_SHOWERR** Do not redirect *stderr* to **/dev/null**.

51022 **WRDE\_UNDEF** Report error on an attempt to expand an undefined shell variable.

51023 The **WRDE\_APPEND** flag can be used to append a new set of words to those generated by a  
 51024 previous call to *wordexp()*. The following rules apply to applications when two or more calls to  
 51025 *wordexp()* are made with the same value of *pwordexp* and without intervening calls to *wordfree()*:

- 51026 1. The first such call shall not set **WRDE\_APPEND**. All subsequent calls shall set it.
- 51027 2. All of the calls shall set **WRDE\_DOOFFS**, or all shall not set it.
- 51028 3. After the second and each subsequent call, *pwordexp->we\_wordv* shall point to a list  
 51029 containing the following:
  - 51030 a. Zero or more null pointers, as specified by **WRDE\_DOOFFS** and *pwordexp->we\_offs*
  - 51031 b. Pointers to the words that were in the *pwordexp->we\_wordv* list before the call, in the  
 51032 same order as before
  - 51033 c. Pointers to the new words generated by the latest call, in the specified order
- 51034 4. The count returned in *pwordexp->we\_wordc* shall be the total number of words from all of  
 51035 the calls.
- 51036 5. The application can change any of the fields after a call to *wordexp()*, but if it does it shall  
 51037 reset them to the original value before a subsequent call, using the same *pwordexp* value, to  
 51038 *wordfree()* or *wordexp()* with the **WRDE\_APPEND** or **WRDE\_REUSE** flag.

51039 If the implementation supports the utilities defined in the Shell and Utilities volume of  
 51040 IEEE Std. 1003.1-200x, and *words* contains an unquoted character—<newline>, ' | ', '&', ';',  
 51041 '<', '>', '(, ' ) ', '{, ' }'—in an inappropriate context, *wordexp()* shall fail, and the number  
 51042 of expanded words shall be 0.

51043 Unless WRDE\_SHOWERR is set in *flags*, *wordexp()* shall redirect *stderr* to */dev/null* for any  
 51044 utilities executed as a result of command substitution while expanding *words*. If  
 51045 WRDE\_SHOWERR is set, *wordexp()* may write messages to *stderr* if syntax errors are detected  
 51046 while expanding *words*.

51047 The application shall ensure that if WRDE\_DOOFFS is set, then *pwordexp->we\_offs* has the same  
 51048 value for each *wordexp()* call and *wordfree()* call using a given *pwordexp*.

51049 The following constants are defined as error return values:

51050 WRDE\_BADCHAR One of the unquoted characters—<newline>, ' | ', '&', ';', '<', '>',  
 51051 ' ( ', ' ) ', ' { ', ' } '—appears in *words* in an inappropriate context.

51052 WRDE\_BADVAL Reference to undefined shell variable when WRDE\_UNDEF is set in *flags*.

51053 WRDE\_CMDSUB Command substitution requested when WRDE\_NOCMD was set in *flags*.

51054 WRDE\_NOSPACE Attempt to allocate memory failed.

51055 WRDE\_SYNTAX Shell syntax error, such as unbalanced parentheses or unterminated  
 51056 string.

#### 51057 RETURN VALUE

51058 Upon successful completion, *wordexp()* shall return 0. Otherwise, a non-zero value, as described  
 51059 in <*wordexp.h*>, shall be returned to indicate an error. If *wordexp()* returns the value  
 51060 WRDE\_NOSPACE, then *pwordexp->we\_wordc* and *pwordexp->we\_wordv* shall be updated to  
 51061 reflect any words that were successfully expanded. In other cases, they shall not be modified.

51062 The *wordfree()* function shall return no value.

#### 51063 ERRORS

51064 No errors are defined.

#### 51065 EXAMPLES

51066 None.

#### 51067 APPLICATION USAGE

51068 The *wordexp()* function is intended to be used by an application that wants to do all of the shell's  
 51069 expansions on a word or words obtained from a user. For example, if the application prompts  
 51070 for a file name (or list of file names) and then uses *wordexp()* to process the input, the user could  
 51071 respond with anything that would be valid as input to the shell.

51072 The WRDE\_NOCMD flag is provided for applications that, for security or other reasons, want to  
 51073 prevent a user from executing shell commands. Disallowing unquoted shell special characters  
 51074 also prevents unwanted side effects, such as executing a command or writing a file.

#### 51075 RATIONALE

51076 This function was included as an alternative to *glob()*. There had been continuing controversy  
 51077 over exactly what features should be included in *glob()*. It is hoped that by providing *wordexp()*  
 51078 (which provides all of the shell word expansions, but which may be slow to execute) and *glob()*  
 51079 (which is faster, but which only performs path name expansion, without tilde or parameter  
 51080 expansion) this will satisfy the majority of applications.

51081 While *wordexp()* could be implemented entirely as a library routine, it is expected that most  
 51082 implementations run a shell in a subprocess to do the expansion.

51083 Two different approaches have been proposed for how the required information might be  
 51084 presented to the shell and the results returned. They are presented here as examples.

51085 One proposal is to extend the *echo* utility by adding a *-q* option. This option would cause *echo* to  
 51086 add a backslash before each backslash and <blank> character that occurs within an argument.

51087 The *wordexp()* function could then invoke the shell as follows:

```
51088 (void) strcpy(buffer, "echo -q");
51089 (void) strcat(buffer, words);
51090 if ((flags & WRDE_SHOWERR) == 0)
51091 (void) strcat(buffer, "2>/dev/null");
51092 f = popen(buffer, "r");
```

51093 The *wordexp()* function would read the resulting output, remove unquoted backslashes, and  
 51094 break into words at unquoted <blank>s. If the WRDE\_NOCMD flag was set, *wordexp()* would  
 51095 have to scan *words* before starting the subshell to make sure that there would be no command  
 51096 substitution. In any case, it would have to scan *words* for unquoted special characters.

51097 Another proposal is to add the following options to *sh*:

51098 **-w wordlist**

51099 This option provides a wordlist expansion service to applications. The words in *wordlist*  
 51100 shall be expanded and the following written to standard output:

- 51101 1. The count of the number of words after expansion, in decimal, followed by a null byte
- 51102 2. The number of bytes needed to represent the expanded words (not including null  
 51103 separators), in decimal, followed by a null byte
- 51104 3. The expanded words, each terminated by a null byte

51105 If an error is encountered during word expansion, *sh* exits with a non-zero status after  
 51106 writing the former to report any words successfully expanded

51107 **-P** Run in “protected” mode. If specified with the **-w** option, no command substitution shall  
 51108 be performed.

51109 With these options, *wordexp()* could be implemented fairly simply by creating a subprocess  
 51110 using *fork()* and executing *sh* using the line:

```
51111 execl(<shell path>, "sh", "-P", "-w", words, (char *)0);
```

51112 after directing standard error to **/dev/null**.

51113 It seemed objectionable for a library routine to write messages to standard error, unless  
 51114 explicitly requested, so *wordexp()* is required to redirect standard error to **/dev/null** to ensure  
 51115 that no messages are generated, even for commands executed for command substitution. The  
 51116 WRDE\_SHOWERR flag can be specified to request that error messages be written.

51117 The WRDE\_REUSE flag allows the implementation to avoid the expense of freeing and  
 51118 reallocating memory, if that is possible. A minimal implementation can call *wordfree()* when  
 51119 WRDE\_REUSE is set.

#### 51120 FUTURE DIRECTIONS

51121 None.

#### 51122 SEE ALSO

51123 *fnmatch()*, *glob()*, the Base Definitions volume of IEEE Std. 1003.1-200x, **<wordexp.h>**, the Shell  
 51124 and Utilities volume of IEEE Std. 1003.1-200x

#### 51125 CHANGE HISTORY

51126 First released in Issue 4. Derived from the ISO POSIX-2 standard.

51127 **Issue 5**

51128 Moved from POSIX2 C-language Binding to BASE.

51129 **Issue 6**

51130 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

51131 The **restrict** keyword is added to the *wordexp()* prototype for alignment with the

51132 ISO/IEC 9899:1999 standard.

51133 **NAME**

51134        **wprintf** — print formatted wide-character output

51135 **SYNOPSIS**

51136        #include <stdio.h>

51137        #include <wchar.h>

51138        int wprintf(const wchar\_t \**format*, ...);

51139 **DESCRIPTION**

51140        Refer to *fwprintf()*.

51141 **NAME**

51142 pwrite, write, writev — write on a file

51143 **SYNOPSIS**

51144 #include &lt;unistd.h&gt;

51145 XSI ssize\_t pwrite(int *fildev*, const void \**buf*, size\_t *nbyte*,  
51146 off\_t *offset*);51147 ssize\_t write(int *fildev*, const void \**buf*, size\_t *nbyte*);

51148 XSI #include &lt;sys/uio.h&gt;

51149 ssize\_t writev(int *fildev*, const struct iovec \**iov*, int *iovcnt*);

51150

51151 **DESCRIPTION**51152 XSI The *pwrite()* function performs the same action as *write()*, except that it writes into a given  
51153 position without changing the file pointer. The first three arguments to *pwrite()* are the same as  
51154 *write()* with the addition of a fourth argument *offset* for the desired position inside the file.51155 **Notes to Reviewers**51156 *This section with side shading will not appear in the final copy. - Ed.*51157 D3, XSH, ERN 676 says that *pwrite()* (and *pread()*) need to be limited to seekable devices or have  
51158 an explicit statement that the *offset* argument is ignored. This item is still open.51159 The *write()* function shall attempt to write *nbyte* bytes from the buffer pointed to by *buf* to the  
51160 file associated with the open file descriptor, *fildev*.51161 If *nbyte* is zero and the file is a regular file, the *write()* function may detect and return errors as  
51162 described below. In the absence of errors, or if error detection is not performed, the *write()*  
51163 function shall return zero and have no other results. If *nbyte* is zero and the file is not a regular  
51164 file, the results are unspecified.51165 On a regular file or other file capable of seeking, the actual writing of data proceeds from the  
51166 position in the file indicated by the file offset associated with *fildev*. Before successful return  
51167 from *write()*, the file offset is incremented by the number of bytes actually written. On a regular  
51168 file, if this incremented file offset is greater than the length of the file, the length of the file shall  
51169 be set to this file offset.51170 On a file not capable of seeking, writing always takes place starting at the current position. The  
51171 value of a file offset associated with such a device is undefined.51172 If the O\_APPEND flag of the file status flags is set, the file offset shall be set to the end of the file  
51173 prior to each write and no intervening file modification operation shall occur between changing  
51174 the file offset and the write operation.51175 XSI If a *write()* requests that more bytes be written than there is room for (for example, the process'  
51176 file size limit or the physical end of a medium), only as many bytes as there is room for shall be  
51177 written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A  
51178 write of 512 bytes shall return 20. The next write of a non-zero number of bytes shall give a  
51179 XSI failure return (except as noted below) and the implementation shall generate a SIGXFSZ signal  
51180 for the thread.51181 If *write()* is interrupted by a signal before it writes any data, it shall return -1 with *errno* set to  
51182 [EINTR].51183 If *write()* is interrupted by a signal after it successfully writes some data, it shall return the  
51184 number of bytes written.

- 51185 If the value of *nbyte* is greater than {SSIZE\_MAX}, the result is implementation-defined.
- 51186 After a *write()* to a regular file has successfully returned:
- 51187 • Any successful *read()* from each byte position in the file that was modified by that write shall
  - 51188 return the data specified by the *write()* for that position until such byte positions are again
  - 51189 modified.
  - 51190 • Any subsequent successful *write()* to the same byte position in the file shall overwrite that
  - 51191 file data.
- 51192 Write requests to a pipe or FIFO shall be handled the same as a regular file with the following
- 51193 exceptions:
- 51194 • There is no file offset associated with a pipe, hence each write request shall append to the
  - 51195 end of the pipe.
  - 51196 • Write requests of {PIPE\_BUF} bytes or less shall not be interleaved with data from other
  - 51197 processes doing writes on the same pipe. Writes of greater than {PIPE\_BUF} bytes may have
  - 51198 data interleaved, on arbitrary boundaries, with writes by other processes, whether or not the
  - 51199 O\_NONBLOCK flag of the file status flags is set.
  - 51200 • If the O\_NONBLOCK flag is clear, a write request may cause the thread to block, but on
  - 51201 normal completion it shall return *nbyte*.
  - 51202 • If the O\_NONBLOCK flag is set, *write()* requests shall be handled differently, in the
  - 51203 following ways:
    - 51204 — The *write()* function shall not block the thread.
    - 51205 — A write request for {PIPE\_BUF} or fewer bytes shall have the following effect: if there is
    - 51206 sufficient space available in the pipe, *write()* shall transfer all the data and return the
    - 51207 number of bytes requested. Otherwise, *write()* shall transfer no data and return  $-1$  with
    - 51208 *errno* set to [EAGAIN].
    - 51209 — A write request for more than {PIPE\_BUF} bytes shall cause one of the following:
      - 51210 — When at least one byte can be written, transfer what it can and return the number of
      - 51211 bytes written. When all data previously written to the pipe is read, it shall transfer at
      - 51212 least {PIPE\_BUF} bytes.
      - 51213 — When no data can be written, transfer no data, and return  $-1$  with *errno* set to
      - 51214 [EAGAIN].
- 51215 When attempting to write to a file descriptor (other than a pipe or FIFO) that supports non-
- 51216 blocking writes and cannot accept the data immediately:
- 51217 • If the O\_NONBLOCK flag is clear, *write()* shall block the calling thread until the data can be
  - 51218 accepted.
  - 51219 • If the O\_NONBLOCK flag is set, *write()* shall not block the thread. If some data can be
  - 51220 written without blocking the thread, *write()* shall write what it can and return the number of
  - 51221 bytes written. Otherwise, it shall return  $-1$  and set *errno* to [EAGAIN].
- 51222 Upon successful completion, where *nbyte* is greater than 0, *write()* shall mark for update the
- 51223 *st\_ctime* and *st\_mtime* fields of the file, and if the file is a regular file, the S\_ISUID and S\_ISGID
- 51224 bits of the file mode may be cleared.
- 51225 XSR If *fildev* refers to a STREAM, the operation of *write()* shall be determined by the values of the
- 51226 minimum and maximum *nbyte* range (packet size) accepted by the STREAM. These values are
- 51227 determined by the topmost STREAM module. If *nbyte* falls within the packet size range, *nbyte*

51228 bytes shall be written. If *nbyte* does not fall within the range and the minimum packet size value  
 51229 is 0, *write()* shall break the buffer into maximum packet size segments prior to sending the data  
 51230 downstream (the last segment may contain less than the maximum packet size). If *nbyte* does not  
 51231 fall within the range and the minimum value is non-zero, *write()* shall fail with *errno* set to  
 51232 [ERANGE]. Writing a zero-length buffer (*nbyte* is 0) to a STREAMS device sends 0 bytes with 0  
 51233 returned. However, writing a zero-length buffer to a STREAMS-based pipe or FIFO sends no  
 51234 message and 0 is returned. The process may issue `I_SWROPT ioctl()` to enable zero-length  
 51235 messages to be sent across the pipe or FIFO.

51236 When writing to a STREAM, data messages are created with a priority band of 0. When writing  
 51237 to a STREAM that is not a pipe or FIFO:

- If `O_NONBLOCK` is clear, and the STREAM cannot accept data (the STREAM write queue is full due to internal flow control conditions), *write()* shall block until data can be accepted.

- If `O_NONBLOCK` is set and the STREAM cannot accept data, *write()* shall return `-1` and set *errno* to [EAGAIN].

- If `O_NONBLOCK` is set and part of the buffer has been written while a condition in which the STREAM cannot accept additional data occurs, *write()* shall terminate and return the number of bytes written.

51245 In addition, *write()* and *writew()* shall fail if the STREAM head has processed an asynchronous  
 51246 error before the call. In this case, the value of *errno* does not reflect the result of *write()* or  
 51247 *writew()*, but reflects the prior error.

51248 XSI The *writew()* function is equivalent to *write()*, but gathers the output data from the *iovcnt* buffers  
 51249 specified by the members of the *iov* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt*-1]. *iovcnt* is valid if  
 51250 greater than 0 and less than or equal to {IOV\_MAX}, defined in <limits.h>.

51251 Each *iovec* entry specifies the base address and length of an area in memory from which data  
 51252 should be written. The *writew()* function shall always write a complete area before proceeding to  
 51253 the next.

51254 If *fildev* refers to a regular file and all of the *iov\_len* members in the array pointed to by *iov* are 0,  
 51255 *writew()* shall return 0 and have no other effect. For other file types, the behavior is unspecified.

51256 If the sum of the *iov\_len* values is greater than {SSIZE\_MAX}, the operation fails and no data is  
 51257 transferred.

51258 SIO If the `O_DSYNC` bit has been set, write I/O operations on the file descriptor complete as defined  
 51259 by synchronized I/O data integrity completion.

51260 If the `O_SYNC` bit has been set, write I/O operations on the file descriptor complete as defined  
 51261 by synchronized I/O file integrity completion.

51262 SHM If *fildev* refers to a shared memory object, the result of the *write()* function is unspecified.

51263 TYM If *fildev* refers to a typed memory object, the result of the *write()* function is unspecified.

51264 For regular files, no data transfer shall occur past the offset maximum established in the open  
 51265 file description associated with *fildev*.

51266 If *fildev* refers to a socket, *write()* is equivalent to *send()* with no flags set.

#### 51267 RETURN VALUE

51268 XSI Upon successful completion, *write()* and *pwrite()* shall return the number of bytes actually  
 51269 written to the file associated with *fildev*. This number shall never be greater than *nbyte*.  
 51270 Otherwise, `-1` shall be returned and *errno* set to indicate the error.

51271 XSI Upon successful completion, `writew()` shall return the number of bytes actually written.  
 51272 Otherwise, it shall return a value of `-1`, the file-pointer shall remain unchanged, and `errno` shall  
 51273 be set to indicate an error.

51274 **ERRORS**

51275 XSI The `write()`, `pwrite()`, and `writew()` functions shall fail if:

51276 [EAGAIN] The `O_NONBLOCK` flag is set for the file descriptor and the thread would be  
 51277 delayed in the `write()` operation.

51278 [EBADF] The `fildev` argument is not a valid file descriptor open for writing.

51279 [EFBIG] An attempt was made to write a file that exceeds the implementation-defined  
 51280 XSI maximum file size or the process' file size limit.

51281 [EFBIG] The file is a regular file, `nbyte` is greater than 0, and the starting position is  
 51282 greater than or equal to the offset maximum established in the open file  
 51283 description associated with `fildev`.

51284 [EINTR] The write operation was terminated due to the receipt of a signal, and no data  
 51285 was transferred.

51286 [EIO] The process is a member of a background process group attempting to write  
 51287 to its controlling terminal, `TOSTOP` is set, the process is neither ignoring nor  
 51288 blocking `SIGTTOU`, and the process group of the process is orphaned. This  
 51289 error may also be returned under implementation-defined conditions.

51290 [ENOSPC] There was no free space remaining on the device containing the file.

51291 [EPIPE] An attempt is made to write to a pipe or FIFO that is not open for reading by  
 51292 any process, or that only has one end open. A `SIGPIPE` signal shall also be sent  
 51293 to the thread.

51294 XSR [ERANGE] The transfer request size was outside the range supported by the `STREAMS`  
 51295 file associated with `fildev`.

51296 MAN The `write()` function shall fail if:

51297 [EAGAIN] or [EWOULDBLOCK]

51298 The file descriptor is for a connection-mode socket, is marked  
 51299 `O_NONBLOCK`, and write would block.

51300 [ECONNRESET] A write was attempted on a connection-mode socket that is not connected.

51301 [EPIPE] A write was attempted on a connection-mode socket that is shut down for  
 51302 writing, or is no longer connected. In the latter case, if the socket is of type  
 51303 `SOCK_STREAM`, the `SIGPIPE` signal is generated to the calling process.  
 51304

51305 The `writew()` function shall fail if:

51306 XSI [EINVAL] The sum of the `iov_len` values in the `iov` array would overflow an `ssize_t`.

51307 XSI The `write()`, `pwrite()`, and `writew()` functions may fail if:

51308 XSR [EINVAL] The `STREAM` or multiplexer referenced by `fildev` is linked (directly or  
 51309 indirectly) downstream from a multiplexer.

51310 MAN [EIO] A physical I/O error has occurred.

51311 MAN [ENOBUFS] Insufficient resources were available in the system to perform the operation.

|       |               |                                                                                                                          |
|-------|---------------|--------------------------------------------------------------------------------------------------------------------------|
| 51312 | [ENXIO]       | A request was made of a nonexistent device, or the request was outside the capabilities of the device.                   |
| 51313 |               |                                                                                                                          |
| 51314 | XSR [ENXIO]   | A hangup occurred on the STREAM being written to.                                                                        |
| 51315 | XSR           | A write to a STREAMS file may fail if an error message has been received at the STREAM head.                             |
| 51316 |               | In this case, <i>errno</i> is set to the value included in the error message.                                            |
| 51317 | MAN           | The <i>write()</i> function may fail if:                                                                                 |
| 51318 | [EACCES]      | A write was attempted on a connection-mode socket and the calling process does not have appropriate privileges.          |
| 51319 |               |                                                                                                                          |
| 51320 | [ENETDOWN]    | A write was attempted on a connection-mode socket and the local network interface used to reach the destination is down. |
| 51321 |               |                                                                                                                          |
| 51322 | [ENETUNREACH] | A write was attempted on a connection-mode socket and no route to the network is present.                                |
| 51323 |               |                                                                                                                          |
| 51324 |               |                                                                                                                          |
| 51325 |               | The <i>writew()</i> function may fail and set <i>errno</i> to:                                                           |
| 51326 | XSI [EINVAL]  | The <i>iovcnt</i> argument was less than or equal to 0, or greater than {IOV_MAX}.                                       |
| 51327 | XSI           | The <i>pwrite()</i> function shall fail and the file pointer remain unchanged if:                                        |
| 51328 | XSI [EINVAL]  | The <i>offset</i> argument is invalid. The value is negative.                                                            |
| 51329 | XSI [ESPIPE]  | <i>fildev</i> is associated with a pipe or FIFO.                                                                         |

51330 **EXAMPLES**51331 **Writing from a Buffer**

51332 The following example writes data from the buffer pointed to by *buf* to the file associated with  
51333 the file descriptor *fd*.

```

51334 #include <sys/types.h>
51335 #include <string.h>
51336 ...
51337 char buf[20];
51338 size_t nbytes;
51339 ssize_t bytes_written;
51340 int fd;
51341 ...
51342 strcpy(buf, "This is a test\n");
51343 nbytes = strlen(buf);
51344 bytes_written = write(fd, buf, nbytes);
51345 ...

```

51346 **Writing Data from an Array**

51347 The following example writes data from the buffers specified by members of the *iov* array to the  
51348 file associated with the file descriptor *fd*.

```
51349 #include <sys/types.h>
51350 #include <sys/uio.h>
51351 #include <unistd.h>
51352 ...
51353 ssize_t bytes_written;
51354 int fd;
51355 char *buf0 = "short string\n";
51356 char *buf1 = "This is a longer string\n";
51357 char *buf2 = "This is the longest string in this example\n";
51358 int iovcnt;
51359 struct iovec iov[3];

51360 iov[0].iov_base = buf0;
51361 iov[0].iov_len = strlen(buf0);
51362 iov[1].iov_base = buf1;
51363 iov[1].iov_len = strlen(buf1);
51364 iov[2].iov_base = buf2;
51365 iov[2].iov_len = strlen(buf2);
51366 ...
51367 iovcnt = sizeof(iov) / sizeof(struct iovec);

51368 bytes_written = writev(fd, iov, iovcnt);
51369 ...
```

51370 **APPLICATION USAGE**

51371 None.

51372 **RATIONALE**

51373 See also the RATIONALE section in *read()*.

51374 An attempt to write to a pipe or FIFO has several major characteristics:

51375 • *Atomic/non-atomic*: A write is atomic if the whole amount written in one operation is not  
51376 interleaved with data from any other process. This is useful when there are multiple writers  
51377 sending data to a single reader. Applications need to know how large a write request can be  
51378 expected to be performed atomically. This maximum is called {PIPE\_BUF}. This volume of  
51379 IEEE Std. 1003.1-200x does not say whether write requests for more than {PIPE\_BUF} bytes  
51380 are atomic, but requires that writes of {PIPE\_BUF} or fewer bytes shall be atomic.

51381 • *Blocking/immediate*: Blocking is only possible with O\_NONBLOCK clear. If there is enough  
51382 space for all the data requested to be written immediately, the implementation should do so.  
51383 Otherwise, the process may block; that is, pause until enough space is available for writing.  
51384 The effective size of a pipe or FIFO (the maximum amount that can be written in one  
51385 operation without blocking) may vary dynamically, depending on the implementation, so it  
51386 is not possible to specify a fixed value for it.

51387 • *Complete/partial/deferred*: A write request:

```
51388 int fildes;
51389 size_t nbyte;
51390 ssize_t ret;
51391 char *buf;
```

51392           ret = write(fildev, buf, nbyte);

51393           may return:

51394           complete    ret=nbyte

51395           partial     ret<nbyte

51396                       This shall never happen if *nbyte*≤{PIPE\_BUF}. If it does happen (with  
51397                       *nbyte*>{PIPE\_BUF}), this volume of IEEE Std. 1003.1-200x does not guarantee  
51398                       atomicity, even if *ret*≤{PIPE\_BUF}, because atomicity is guaranteed according  
51399                       to the amount *requested*, not the amount *written*.

51400           deferred:   ret=-1, *errno*=[EAGAIN]

51401                       This error indicates that a later request may succeed. It does not indicate that it  
51402                       *shall* succeed, even if *nbyte*≤{PIPE\_BUF}, because if no process reads from the  
51403                       pipe or FIFO, the write never succeeds. An application could usefully count the  
51404                       number of times [EAGAIN] is caused by a particular value of  
51405                       *nbyte*>{PIPE\_BUF} and perhaps do later writes with a smaller value, on the  
51406                       assumption that the effective size of the pipe may have decreased.

51407                       Partial and deferred writes are only possible with O\_NONBLOCK set.

51408           The relations of these properties are shown in the following tables:

51409

51410

51411

51412

51413

51414

| Write to a Pipe or FIFO with O_NONBLOCK clear |                                 |                                 |                                  |
|-----------------------------------------------|---------------------------------|---------------------------------|----------------------------------|
| Immediately Writable:                         | None                            | Some                            | <i>nbyte</i>                     |
| <i>nbyte</i> ≤{PIPE_BUF}                      | Atomic blocking<br><i>nbyte</i> | Atomic blocking<br><i>nbyte</i> | Atomic immediate<br><i>nbyte</i> |
| <i>nbyte</i> >{PIPE_BUF}                      | Blocking <i>nbyte</i>           | Blocking <i>nbyte</i>           | Blocking <i>nbyte</i>            |

51415           If the O\_NONBLOCK flag is clear, a write request shall block if the amount writable  
51416           immediately is less than that requested. If the flag is set (by *fcntl()*), a write request shall never  
51417           block.

51418

51419

51420

51421

51422

51423

| Write to a Pipe or FIFO with O_NONBLOCK set |              |                                   |                                   |
|---------------------------------------------|--------------|-----------------------------------|-----------------------------------|
| Immediately Writable:                       | None         | Some                              | <i>nbyte</i>                      |
| <i>nbyte</i> ≤{PIPE_BUF}                    | -1, [EAGAIN] | -1, [EAGAIN]                      | Atomic <i>nbyte</i>               |
| <i>nbyte</i> >{PIPE_BUF}                    | -1, [EAGAIN] | < <i>nbyte</i> or -1,<br>[EAGAIN] | ≤ <i>nbyte</i> or -1,<br>[EAGAIN] |

51424           There is no exception regarding partial writes when O\_NONBLOCK is set. With the exception  
51425           of writing to an empty pipe, this volume of IEEE Std. 1003.1-200x does not specify exactly when  
51426           a partial write is performed since that would require specifying internal details of the  
51427           implementation. Every application should be prepared to handle partial writes when  
51428           O\_NONBLOCK is set and the requested amount is greater than {PIPE\_BUF}, just as every  
51429           application should be prepared to handle partial writes on other kinds of file descriptors.

51430           The intent of forcing writing at least one byte if any can be written is to assure that each write  
51431           makes progress if there is any room in the pipe. If the pipe is empty, {PIPE\_BUF} bytes must be  
51432           written; if not, at least some progress must have been made.

51433           Where this volume of IEEE Std. 1003.1-200x requires -1 to be returned and *errno* set to  
51434           [EAGAIN], most historical implementations return zero (with the O\_NDELAY flag set, which is  
51435           the historical predecessor of O\_NONBLOCK, but is not itself in this volume of

- 51436 IEEE Std. 1003.1-200x). The error indications in this volume of IEEE Std. 1003.1-200x were chosen  
 51437 so that an application can distinguish these cases from end-of-file. While *write()* cannot receive  
 51438 an indication of end-of-file, *read()* can, and the two functions have similar return values. Also,  
 51439 some existing systems (for example, Eighth Edition) permit a write of zero bytes to mean that  
 51440 the reader should get an end-of-file indication; for those systems, a return value of zero from  
 51441 *write()* indicates a successful write of an end-of-file indication.
- 51442 Implementations are allowed, but not required, to perform error checking for *write()* requests of  
 51443 zero bytes.
- 51444 The concept of a {PIPE\_MAX} limit (indicating the maximum number of bytes that can be  
 51445 written to a pipe in a single operation) was considered, but rejected, because this concept would  
 51446 unnecessarily limit application writing.
- 51447 See also the discussion of O\_NONBLOCK in *read()*.
- 51448 Writes can be serialized with respect to other reads and writes. If a *read()* of file data can be  
 51449 proven (by any means) to occur after a *write()* of the data, it must reflect that *write()*, even if the  
 51450 calls are made by different processes. A similar requirement applies to multiple write operations  
 51451 to the same file position. This is needed to guarantee the propagation of data from *write()* calls  
 51452 to subsequent *read()* calls. This requirement is particularly significant for networked file  
 51453 systems, where some caching schemes violate these semantics.
- 51454 Note that this is specified in terms of *read()* and *write()*. Additional calls, such as the common  
 51455 *readv()* and *writew()*, would want to obey these semantics. A new “high-performance” write  
 51456 analog that did not follow these serialization requirements would also be permitted by this  
 51457 wording. This volume of IEEE Std. 1003.1-200x is also silent about any effects of application-  
 51458 level caching (such as that done by *stdio*).
- 51459 This volume of IEEE Std. 1003.1-200x does not specify the value of the file offset after an error is  
 51460 returned; there are too many cases. For programming errors, such as [EBADF], the concept is  
 51461 meaningless since no file is involved. For errors that are detected immediately, such as  
 51462 [EAGAIN], clearly the pointer should not change. After an interrupt or hardware error, however,  
 51463 an updated value would be very useful and is the behavior of many implementations.
- 51464 This volume of IEEE Std. 1003.1-200x does not specify behavior of concurrent writes to a file  
 51465 from multiple processes. Applications should use some form of concurrency control.
- 51466 **FUTURE DIRECTIONS**
- 51467 None.
- 51468 **SEE ALSO**
- 51469 *chmod()*, *creat()*, *dup()*, *fcntl()*, *getrlimit()*, *lseek()*, *open()*, *pipe()*, *ulimit()*, the Base Definitions  
 51470 volume of IEEE Std. 1003.1-200x, <limits.h>, <stropts.h>, <sys/uio.h>, <unistd.h>
- 51471 **CHANGE HISTORY**
- 51472 First released in Issue 1. Derived from Issue 1 of the SVID.
- 51473 **Issue 4**
- 51474 The <unistd.h> header is added to the SYNOPSIS section.
- 51475 Reference to *ulimit* in the DESCRIPTION is marked as an extension.
- 51476 Reference to the process' file size limit and the *ulimit()* function are marked as extensions in the  
 51477 description of the [EFBIG] error.
- 51478 The [ENXIO] error is marked as an extension.
- 51479 The APPLICATION USAGE section is removed.

51480 The description of [EINTR] is amended.

51481 The following changes are incorporated for alignment with the ISO POSIX-1 standard:

- 51482 • The type of the argument *buf* is changed from **char\*** to **const void\***, and the type of the
- 51483 argument *nbyte* is changed from **unsigned** to **size\_t**.
- 51484 • The DESCRIPTION is changed as follows:
  - 51485 — Writing at end-of-file is atomic.
  - 51486 — {SSIZE\_MAX} is now used to determine the maximum value of *nbyte*.
  - 51487 — The consequences of activities after a call to the *write()* function are added.
  - 51488 — To improve clarity, the text describing operations on pipes or FIFOs when
  - 51489 O\_NONBLOCK is set is restructured.

#### 51490 **Issue 4, Version 2**

51491 The following changes are incorporated for X/OPEN UNIX conformance:

- 51492 • The *writew()* function is added to the SYNOPSIS.
- 51493 • The DESCRIPTION is updated to describe the writing of data to STREAMS files, an
- 51494 operational description of the *writew()* function is included, and a statement is added
- 51495 indicating that SIGXFSZ is generated if an attempted write operation would cause the
- 51496 maximum file size to be exceeded.
- 51497 • The RETURN VALUE section is updated to describe values returned by the *writew()* function.
- 51498 • The ERRORS section has been restructured to describe errors that apply to both *write()* and
- 51499 *writew()* apart from those that apply to *writew()* specifically. The [EIO], [ERANGE], and
- 51500 [EINVAL] errors are also added.

#### 51501 **Issue 5**

51502 The DESCRIPTION is updated for alignment with the POSIX Realtime Extension and the POSIX

51503 Threads Extension.

51504 Large File Summit extensions are added.

51505 The *pwrite()* function is added.

#### 51506 **Issue 6**

51507 The DESCRIPTION states that the *write()* function does not block the thread. Previously this

51508 said “process” rather than “thread”.

51509 The DESCRIPTION and ERRORS sections are updated so that references to STREAMS are

51510 marked as part of the XSI STREAMS Option Group.

51511 The following new requirements on POSIX implementations derive from alignment with the

51512 Single UNIX Specification:

- 51513 • The DESCRIPTION now states that if *write()* is interrupted by a signal after it has
- 51514 successfully written some data, it returns the number of bytes written. In earlier versions of
- 51515 this volume of IEEE Std. 1003.1-200x, it was optional whether *write()* returned the number of
- 51516 bytes written, or whether it returned  $-1$  with *errno* set to [EINTR]. This is a FIPS requirement.
- 51517 • The following changes are made to support large files:
  - 51518 — For regular files, no data transfer occurs past the offset maximum established in the open
  - 51519 file description associated with the *files*.

- 51520 — A second [EFBIG] error condition is added.
- 51521 • The [EIO] error condition is added.
- 51522 • The [EPIPE] error condition is added for when a pipe has only one end open.
- 51523 • The [ENXIO] optional error condition is added.
- 51524 Text referring to sockets is added to the DESCRIPTION.
- 51525 The following changes were made to align with the IEEE P1003.1a draft standard:
- 51526 • The effect of reading zero bytes is clarified.
- 51527 The DESCRIPTION is updated for alignment with IEEE Std. 1003.1j-2000 by specifying that  
51528 *write()* results are unspecified for typed memory objects.
- 51529 The following error conditions are added for operations on sockets: [EAGAIN],  
51530 [EWOULDBLOCK], [ECONNRESET], [ENOTCONN], and [EPIPE].
- 51531 The [EIO] error is changed to “may fail”.
- 51532 The [ENOBUFS] error is added for sockets.
- 51533 The following error conditions are added for operations on sockets: [EACCES], [ENETDOWN],  
51534 and [ENETUNREACH].

51535 **NAME**

51536           wscanf — convert formatted wide-character input

51537 **SYNOPSIS**

51538           #include &lt;stdio.h&gt;

51539           #include &lt;wchar.h&gt;

51540           int wscanf(const wchar\_t \**format*, ... );51541 **DESCRIPTION**51542           Refer to *fwscanf()*.

51543 **NAME**

51544 y0, y1, yn — Bessel functions of the second kind

51545 **SYNOPSIS**

```
51546 xSI #include <math.h>
51547 double y0(double x);
51548 double y1(double x);
51549 double yn(int n, double x);
51550
```

51551 **DESCRIPTION**

51552 The *y0()*, *y1()*, and *yn()* functions shall compute Bessel functions of *x* of the second kind of  
 51553 orders 0, 1, and *n* respectively. The application shall ensure that the value of *x* is positive.

51554 An application wishing to check for error situations should set *errno* to 0 before calling *y0()*,  
 51555 *y1()*, or *yn()*. If *errno* is non-zero on return, or the return value is NaN, an error has occurred.

51556 **RETURN VALUE**

51557 Upon successful completion, *y0()*, *y1()*, and *yn()* shall return the relevant Bessel value of *x* of  
 51558 the second kind.

51559 If *x* is NaN, NaN shall be returned and *errno* may be set to [EDOM].

51560 If the *x* argument to *y0()*, *y1()*, or *yn()* is negative, *-HUGE\_VAL* or NaN shall be returned, and  
 51561 *errno* may be set to [EDOM].

51562 If *x* is 0.0, *-HUGE\_VAL* shall be returned and *errno* may be set to [ERANGE] or [EDOM].

51563 If the correct result would cause underflow, 0.0 shall be returned and *errno* may be set to  
 51564 [ERANGE].

51565 If the correct result would cause overflow, *-HUGE\_VAL* or 0.0 shall be returned and *errno* may  
 51566 be set to [ERANGE].

51567 **ERRORS**

51568 The *y0()*, *y1()*, and *yn()* functions may fail if:

51569 [EDOM] The value of *x* is negative or NaN. |

51570 [ERANGE] The value of *x* is too large in magnitude, or *x* is 0.0, or the correct result would |  
 51571 cause overflow or underflow.

51572 No other errors shall occur.

51573 **EXAMPLES**

51574 None.

51575 **APPLICATION USAGE**

51576 None.

51577 **RATIONALE**

51578 None.

51579 **FUTURE DIRECTIONS**

51580 None.

51581 **SEE ALSO**

51582 *isnan()*, *j0()*, the Base Definitions volume of IEEE Std. 1003.1-200x, <math.h> |

51583 **CHANGE HISTORY**

51584 First released in Issue 1. Derived from Issue 1 of the SVID.

51585 **Issue 4**

51586 References to *matherr()* are removed.

51587 The RETURN VALUE and ERRORS sections are substantially rewritten to rationalize error  
51588 handling in the mathematics functions.

51589 **Issue 5**

51590 The DESCRIPTION is updated to indicate how an application should check for an error. This  
51591 text was previously published in the APPLICATION USAGE section.

51592 **Issue 6**

51593 The DESCRIPTION is updated to avoid use of the term “must” for application requirements.

